Measuring Independence of Datasets

Vladimir Braverman Rafail Ostrovsky
University of California Los Angeles
{vova, rafail}@cs.ucla.edu

March 1, 2009

Abstract

Measuring independence between two or more random variables is a fundamental problem that touches many areas of computer science. The problems of efficiently testing pairwise, or k-wise, independence were recently considered by Alon, Andoni, Kaufman, Matulef, Rubinfeld and Xie (STOC 07); Alon, Goldreich and Mansour (IPL 03); Batu, Fortnow, Fischer, Kumar, Rubinfeld and White (FOCS 01); and Batu, Kumar and Rubinfeld (STOC 04). They addressed the problem of minimizing the number of samples needed to obtain sufficient approximation, when the joint distribution is accessible through a sampling procedure.

A data stream model represents another setting where approximating pairwise, or k-wise, independence with sublinear memory is of considerable importance. Unlike the work in the aforementioned papers, in the streaming model the joint distribution is given by a stream of k-tuples, with the goal of testing correlations among the components measured over the entire stream. In the streaming model, Indyk and McGregor (SODA 08) recently gave exciting new results for measuring pairwise independence.

Statistical distance is one of the most fundamental metrics for measuring the similarity of two distributions, and it has been a metric of choice in many papers that discuss distribution closeness (see, for example, Rubinfeld and Servedio (STOC 05); Sahai and Vadhan (JACM 03); and the above papers). The Indyk and McGregor methods provide $\log n$-approximation under statistical distance between the joint and product distributions in the streaming model. (In contrast, for the L_2 metric, Indyk and McGregor give an $(1 \pm \epsilon)$-approximation for the same problem, but for probability distributions, statistical distance is a significantly more powerful metric then the L_2 metric). For the L_1 metric, in addition to $\log n$ approximation, Indyk and McGregor give an ϵ-approximation that requires linear memory, and also give a method that requires two passes to solve a promise problem for a restricted range of parameters. Indyk and McGregor leave, as their main open question, the problem of improving their $\log n$-approximation for the statistical distance metric.

In this paper we solve the main open problem posed by of Indyk and McGregor for the statistical distance for pairwise independence and extend this result to any constant k. In particular, we present an algorithm that computes an (ϵ, δ)-approximation of the statistical distance between the joint and product distributions defined by a stream of k-tuples. Our algorithm requires $O(\left(\frac{1}{\epsilon}\log(\frac{nm}{\delta})\right)^{(30+k)^k})$ memory and a single pass over the data stream.

1 Introduction

Finding correlations between columns of a table is a fundamental problem in databases. Virtually all commercial databases construct query plans for queries that employ cross-dimensional predicates. The basic step is estimating "selectivity" (i.e., the number of rows that satisfy the predicate conditions) of the complex predicate. Without any prior knowledge, the typical solution is to compute selectivity of each column separately and use the multiplication as an estimate. Thus, optimizers make a "statistical independence assumption" which sometimes may not hold. Incorrect estimations may lead to suboptimal query plans and decrease performance significantly. Identifying correlations between database columns by measuring a level of independence between columns has a long history in the database research community. To illustrate this point, we cite as an example, Poosala and Ioannidis [39]:

> *"For a query involving two or more attributes of the same relation, its result size depends on the joint data distribution of those attributes; i.e., the frequencies of all combinations of attribute values in the database. Due to the multi-dimensional nature of these distributions and the large number of such attribute value combinations, direct approximation of joint distributions can be rather complex and expensive. In practice, most commercial DBMSs adopt the attribute value independence assumption. Under this assumption, the data distributions of individual attributes in a relation are independent of each other and the joint data distribution can be derived from the individual distributions (which are approximated by one-dimensional histograms). Unfortunately, real-life data rarely satisfies the attribute value independence assumption. For instance, functional dependencies represent the exact opposite of the assumption. Moreover, there are intermediate situations as well. For example, it is natural for the salary attribute of the Employee relation to be strongly dependent on the age attribute (i.e., higher/lower salaries mostly going to older/younger people). Making the attribute value independence assumption in these cases may result in very inaccurate approximations of joint data distributions and therefore inaccurate query result size estimations with devastating effects on a DBMS's performance."*

For data warehouses, it is important to find correlated columns for correct schema construction, as Kimball and Caserta note in [34]:

> "*Perfectly correlated attributes, such as the levels of a hierarchy, as well as attributes with a reasonable statistical correlation, should be part of the same dimension.*"

In practice, typical solutions for finding correlations between columns are either histograms (see e.g., [39]) or sampling (see e.g., [30]). These methods have their natural disadvantages, i.e., they do not tolerate deletions and may require several passes over the data. When it comes to very large data volumes, it is critical to maintain sublinear in terms of memory solutions that do not require additional passes over the data and can tolerate incremental updates of the data, e.g., deletions.

For these purposes, a theoretical *data stream model* can be useful. For data warehouses, the "loading" phase of the ETL process (see e.g., Kimball and Caserta [34]) can be seen as a data stream. When reading a database table, the process can be considered as a stream of data tuples. Thus, the data stream model represents another setting where approximating pairwise or k-wise independence with sublinear memory is of considerable importance.

1.1 Precise Definition of the Problem

The natural way to model database tables in a streaming model is by considering a stream of tuples. In this paper we consider a stream of k-tuples $(i_1, \ldots, i_k)$ where $i_l \in [n]$. (For simplicity, we assume that elements of all columns are drawn from the same domain, even though our approach trivially extends to a general case of different domains.) As pointed out in [39, 30, 32], the natural way to define a joint distribution of two (or more) columns is given by the frequencies of all combinations of coordinates. Similarly, the distribution of each column is defined by the corresponding set of frequencies; the definition of a product distribution follows. Let us define these notions precisely.[1]

Definition 1.1. *Let D be a stream of elements $p_1, \ldots, p_m$, where each stream element is a k-tuple $\mathbf{i} = (\mathbf{i}_1, \ldots, \mathbf{i}_k)$, where $\mathbf{i}_l \in [n]$. A* frequency *of a tuple $\mathbf{i} \in [n]^k$ is defined as the number of times it appears in*

[1] Here and thenceforth, we use lowercase Latin characters for indexes. We use an italic font for integers and a boldface font for multidimensional indexes, e.g., $i \in [n]$ and $\mathbf{i} \in [n]^k$. For a multidimensional index, we use subscript to indicate its coordinate, e.g., $\mathbf{i}_1$ indicates the first coordinate of $\mathbf{i}$.

D: $f_{\mathbf{i}} = |\{j : p_j = \mathbf{i}\}|$. *For* $l \in [k]$, *a* l-th margin frequency *of* $t \in [n]$ *is the number of times* t *appears as a* l*-th coordinate:* $f_l(t) = \sum_{\mathbf{i} \in [n]^k, \mathbf{i}_l = t}^{n} f_{\mathbf{i}}$. *A* joint distribution *is defined by a vector of probabilities* $P_{joint}(\mathbf{i}) = \frac{f_{\mathbf{i}}}{m}, \mathbf{i} \in [n]^k$. *Here* m *is the size of stream* D. *A* l-th margin distribution *is defined by a vector of probabilities* $P_l(t) = \frac{f_l(t)}{m}$, $t \in [n]$. *A* product distribution *is defined as:* $P_{product}(\mathbf{i}) = \prod_{l=1}^{k} P_l(\mathbf{i}_l), \mathbf{i} \in [n]^k$.

Statistical distance is one of the most fundamental metrics for measuring the similarity of two distributions, and it has been a metric of choice in many papers that discuss distribution closeness (see e.g., [2, 4, 10, 12, 32, 41, 40]). Given two distributions over a discrete domain, the statistical distance is half of L_1 distance between the probability vectors.

Definition 1.2. *Consider two distributions over a finite domain* Ω *given by two random variables* V, U. Statistical distance $\Delta(V, U)$ *is defined as:*

$$\Delta(V, U) = \frac{1}{2} \sum_{x \in \Omega} |P(V = x) - P(U = x)| = \max_{B \subseteq \Omega} |P(V \in B) - P(U \in B)|.$$

In particular, one of the most common methods of measuring independence is computing statistical distance between product and joint distributions (see e.g., [10, 32]). This is precisely the way we define our problem:

Definition 1.3. *An* Independence Problem *is the following: Given stream* D *of* k*-tuples, approximate, with one pass over* D*, with small memory and high precision the statistical distance between joint and product distribution* $\Delta(P_{joint}, P_{product})$.

In the streaming model, Indyk and McGregor [32] recently gave exciting new results for measuring pairwise independence, i.e., for $k = 2$. To measure the independence, they consider two metrics: L_2 and L_1. Recall that the L_2 distance between two probability distributions is a L_2 distance of their probability vectors. In particular, the independence problem under the L_2 metric is defined as $\|P_{joint} - P_{product}\|_2$.

For the L_2 metric and $k = 2$, Indyk and McGregor give an $(1 \pm \epsilon)$-approximation using polylogarithmic space. However, it is well known that for probability distributions, statistical distance is a significantly more powerful metric then the L_2 metric. For instance, consider two distributions on $[2n]$, where the first distribution is uniform on $\{1, \ldots, n\}$ and the second is uniform on $\{n+1, \ldots, 2n\}$. In this case the statistical

distance is 1 but the L_2 distance is $\sqrt{2/n} \to 0$. For example, Batu, Fortnow, Rubinfeld, Smith and White [11] say:

> *"However, the L_2-distance does not in general give a good measure of the closeness of two distributions. For example, two distributions can have disjoint support and still have small L_2-distance."*

For the statistical distance metric and $k = 2$, the Indyk and McGregor methods provide $\log n$-approximation with polylogarithmic memory. In addition to $\log n$-approximation, Indyk and McGregor give an $(1 \pm \epsilon)$-approximation that requires $\Omega(n)$ memory, and also give a method that requires two passes to solve a promise problem for a restricted range of parameters. Indyk and McGregor leave, as their main open question, the problem of improving their $\log n$-approximation for the statistical distance metric.

In this paper we solve the main open problem posed by of Indyk and McGregor for the statistical distance for pairwise independence and extend this result to any constant k. In particular, we present an algorithm that computes an (ϵ, δ)-approximation of the statistical distance between the joint and product distributions defined by a stream of k-tuples. Our algorithm requires $O((\frac{1}{\epsilon}\log(\frac{nm}{\delta}))^{(30+k)^k})$ memory and a single pass over the data stream. Theorem 2.5 formally describes our main result. We did not try to optimize the constants in our memory bounds.

1.2 Implicit Tensors

It is convenient to present an alternative, equivalent formulation of the *independence problem* as well. We can consider the problem of approximating the sum of absolute values of a *tensor* M_{Ind}.

Definition 1.4. *An s-dimensional* tensor M *is a s-dimensional array with indexes in the range $[n]$; that is, M has an entry for each $\mathbf{i} \in [n]^s$. We denote by $m_{\mathbf{i}}$ the $\mathbf{i}$-th entry of M for each $\mathbf{i} \in [n]^s$.*

Definition 1.5. *Let M be a s-dimensional tensor with entries $m_{\mathbf{i}}, \mathbf{i} \in [n]^s$. An L_1-norm of M is a $|M| = \sum_{\mathbf{i}\in[n]^s} |m_{\mathbf{i}}|$.*

For example, a 1-dimensional tensor is an n-dimensional vector, a 2-dimensional tensor is an $n \times n$-matrix and so forth.

Many streaming problems address *explicitly* defined vectors (or matrices) where entries are equal to frequencies of corresponding stream elements. The Independence problem diverges from this setting; e.g., for pairwise independence, a pair (i, j) affects all entries in i-th row and j-th column of the product probability matrix. To reflect this important difference we consider the case where the entries of a tensor are defined *implicitly* by a data stream.

Definition 1.6. *Let $\mathcal{D}$ be a collection of data streams of size m of elements from domain Ω. Let $\mathcal{F} : \mathcal{D} \times [n]^s \mapsto R$ be a fixed function. We say that s-dimensional tensor M with entries $m_{\mathbf{i}} = \mathcal{F}(D, \mathbf{i}), \mathbf{i} \in [n]^s$ is* implicity defined *by $\mathcal{F}$, given D. We denote an implicitly defined tensor as $\mathcal{F}(D)$.*

Definition 1.7. *Let $\mathcal{D}$ be a collection of data streams of size m of k-tuples from domain $[n]^k$. A k-wise* Independence Function *$\mathcal{F}_{Ind} : \mathcal{D} \times [n]^k \mapsto R$ is a function defined as $\mathcal{F}_{Ind}(D, \mathbf{i}) = m^k f_{\mathbf{i}} - \prod_{l=1}^{k} f_l(\mathbf{i}_l)$ for $\mathbf{i} \in [n]^k$. Here $f_{\mathbf{i}}$ is given by Definition 1.1.* Statistical distance tensor *M_{Ind} is a k-dimensional tensor implicitly defined by $\mathcal{F}_{Ind}$, i.e., $M_{Ind} = \mathcal{F}_{Ind}(D)$.*

The main objective of our paper is approximating $|M_{Ind}|$. In particular, this implies solving the Independence problem since $\Delta(P_{joint}, P_{product}) = \frac{1}{m^k}|M_{Ind}|$, and since $m = |D|$ can be computed precisely. We thus freely interchange the notions of the independence problem and computing $|M_{Ind}|$. In fact, our approach is applicable to any function $\mathcal{F}$ for which conditions of our main theorems are true.

1.3 Why Existing Methods for Estimating L_1 Do Not Work

Alon, Matias and Szegedy [5] initiated the study of computing norms of vectors defined by a data stream. In their setting vector entries are defined by frequencies of the corresponding elements in the stream. Their influential paper was followed by a sequence of exciting results including, among many others, works by Bhuvanagiri, Ganguly, Kesh and Saha [14]; Charikar, Chen and Farach-Colton [17]; Cormode and Muthukrishnan [21, 22]; Feigenbaum, Kannan, Strauss and Viswanathan [26]; Ganguly and Cormode [29]; Indyk [31]; Indyk and Woodruff [33]; and Li [35, 36].

There is an important difference between settings of [5] and the Independence problem. Indeed, while the entries of the independence tensor are defined by frequencies of tuples, there is no linear dependence. As a result, the aforementioned algorithms are not directly applicable to the Independence problem.

To illustrate this point, consider the celebrated method of stable distributions by Indyk [31]. For L_1 norm, Indyk observed that a polylogarithmic (in terms of n and m) number of sketches of the form $\sum C_i v_i$ gives an $(1 \pm \epsilon)$-approximation of $|V|$, when C_i are independent random variables with Cauchy distribution. Let us discuss the applicability of this method to the problem of pairwise independence. A sketch $\sum_{\mathbf{i}} C_{\mathbf{i}} m_{\mathbf{i}}$, $\mathbf{i} \in [n]^2$, would solve this problem; unfortunately, it is not clear how to construct a sketch in this form. In particular, the probability matrix of the product distribution is given implicitly as two vectors of margin sketches. It is not hard to construct sketches for margin distributions; however, it is not at all clear how to obtain a sketch for product distribution without using a multiplication of margin sketches. On the other hand, if we do use a multiplication of margin sketches (this is the approach of Indyk and McGregor), the random variable that is associated with the tensor's elements is a product of independent Cauchy variables. Therefore, random variables for distinct entries are *not independent*, and thus typical arguments used for stable distribution methods do not work anymore. In fact, the main focus of the Indyk and McGregor analysis is to overcome this problem:

> "*Perhaps ironically, the biggest technical challenges that arise relate to ensuring that different components of our estimates are sufficiently independent.*"

For pairwise independence, Indyk and McGregor use the product of two Cauchy variables, where one of them is "truncated." Using elegant observations, they show that such a sketch allows achieving $\log n$-approximation of the statistical distance. Unfortunately, it is not clear how the method of a Cauchy product can be improved at all, since the $\log n$ factor is a necessary component of their seemingly tight analysis.

1.4 A Description of Our Approach

As we discuss below, solving the Independence problem requires developing multiple new tools and using them jointly with known methods.

Dimension Reduction for Implicit Tensors. Our solution can be logically divided into three steps which are explained, informally, below.

First, we prove that given a $polylog$-approximation algorithm for k-dimensional tensors and an ϵ-approximation algorithm for a special type of $(k-1)$-dimensional tensors, it is possible to derive an ϵ-

approximation algorithm on k-dimensional tensors, where the resulting algorithm increases memory bound by a factor $O((\frac{1}{\epsilon}\log\frac{nm}{\delta})^{O(1)})$. Thus, we can trade dimensionality and precision for memory. To illustrate this step, consider pairwise independence. There exist an ϵ-approximation algorithm on vectors [31] and a $\log n$-approximation algorithm on matrices [32]. We show that these algorithms can be used to obtain an ϵ-approximation algorithm on matrices. This informal idea is stated precisely as Dimension Reduction Theorem 2.1. This theorem is the main technical contribution of our paper; the majority of the paper is devoted to establishing its validity.

Second, given a $polylog$-approximation algorithm for k-dimensional tensors and an ϵ-approximation algorithm on vectors, we can derive an ϵ-approximation algorithm on k-dimensional tensors by applying the Dimension Reduction Theorem recursively k-times. The memory will be increased by a factor roughly $O((\frac{1}{\epsilon}\log\frac{nm}{\delta})^{(30+k)^k})$ which is $O((\frac{1}{\epsilon}\log\frac{nm}{\delta})^{O(1)})$ for constant k. This informal idea is stated precisely as Theorem 2.2.

Third, we show that the conditions for Theorem 2.2 hold for the Independence problem. These results are stated in Lemmas 2.4 and 2.3, and in fact are a generalization of results from [31, 32]. Section 6 is devoted to the proof of these lemmas.

The rest of our discussion is devoted to a description of the main ideas behind the Dimension Reduction Theorem.

Hyperplanes and Absolute Vectors. Consider a matrix M; a very natural idea to approximate $|M|$ is by approximating a L_1 norm of a vector with entries equal to L_1 norms of rows of M. We generalize this idea to tensors by defining the following operators.

Definition 1.8. *For any $s,t \geq 0$, we denote by $(,)$ a mapping from $[n]^s \times [n]^t$ to $[n]^{s+t}$ obtained by concatenation of coordinates. For instance, $((1,2),3)$ is a an element from $[n]^3$ with coordinates $1,2,3$ respectfully.*

Definition 1.9. *Let M be a s-dimensional tensor with entries $m_{\mathbf{j}}, \mathbf{j} \in [n]^s$. For any $l \in [n]$, $Hyperplane(M,l)$ is a $(s-1)$-dimensional tensor with entries $m_{(l,\mathbf{i})}$ for $\mathbf{i} \in [n]^{s-1}$.*

For example, when $k=2$, the l-th hyperplane of a matrix M is its l-th row.

Definition 1.10. *An l-th hyperplane is* α-significant *if* $|Hyperplane(M,l)| \geq \alpha|M|$.

For example, when $k = 2$, the l-th row is α-significant if the L_1-norm of the vector defined by the l-th row carries at least α-fraction of $|M|$.

Definition 1.11. *For a s-dimensional tensor M, an $AbsoluteVector(M)$ is a vector of dimensionality n with entries $|Hyperplane(M,l)|, l \in [n]$. In particular, $|AbsoluteVector(M)| = |M|$.*

Projected Dimensions. To prove Dimension Reduction Theorem 2.1 we need to map s-dimensional tensors to $(s-1)$-dimensional tensors with a small distortion of L_1. We come up with the following mapping.

Definition 1.12. *Let M be a s-dimensional tensor with entries $m_{\mathbf{l}}$, where $\mathbf{l} \in [n]^s$, and let $0 \leq t \leq s$. A* Suffix-Sum *tensor $T_t(M)$ is a $(s-t)$-dimensional tensor with entries (for each $\mathbf{i} \in [n]^{s-t}$):*

$$m'_{\mathbf{i}} = \sum_{\mathbf{j} \in [n]^t} m_{(\mathbf{j},\mathbf{i})}$$

Also, we define $T_0(M) = M$. In other words, the $\mathbf{i}$-th entry of $T_t(M)$ is obtained by summing all elements of M with the $(s-t)$-suffix equal to $\mathbf{i}$. In particular, $T_s(M)$ is a scalar that is equal to $\sum_{\mathbf{i} \in [n]^s} m_{\mathbf{i}}$.

For matrix M with entries $m_{i,j}$, the Suffix-Sum operator $T_1(M)$ defines a vector V with entries $v_j = \sum_i m_{i,j}$. In other words, all entries of M that belong to the same columns (i.e., have the same second coordinate, i.e., the same "suffix") are "summed-up" to generate a single entry of V. In some sense, the Suffix-Sum operator is orthogonal to the AbsoluteVector operator. In the latter case we sum up the absolute values that belong to the same hyperplane, i.e., have identical prefix; in the former case we sum up all elements (and not their absolute values) that have an identical suffix.

Clearly $|T_1(M)| \leq |M|$; however, it is possible in general that $|T_1(M)| \ll |M|$. The key observation is that in some cases $|T_1(M)| \sim |M|$ and thus we can use an approximation of $|T_1(M)|$ to approximate $|M|$. To illustrate this point, consider a matrix M with entries $m_{i,j}$ that contains a very "significant" row i (i.e., $\sum_j |m_{i,j}| \sim |M|$). The key observation is that in this case $|T_1(M)| \sim |M|$; thus, if there is a significant row, it can approximated using $|T_1(M)|$. The same idea is easily generalized: if a s-dimensional tensor M contains a $(1-\epsilon)$-significant hyperplane $Hyperplane(M,l)$, then $|T_1(M)|$ is an 2ϵ-approximation of $|Hyperplane(M,l)|$. We prove this statement in Fact 3.6.

Note that $T_1(M)$ is a $(s-1)$-dimensional tensor; if M is a matrix, then $T_1(M)$ is a vector for which we can apply methods from [31]. Thus, approximating $|T_1(M)|$ is potentially an easier problem.

Certifying Tournaments. We have shown that $T_1(M)$ can be useful for approximating $|M|$. However, when can we rely on the value of $|T_1(M)|$? In particular, how can we distinguish between the cases when there is a heavy hyperplane (and thus $|T_1(M)|$ is a good approximation) and the case when there is no heavy hyperplane (and thus $|T_1(M)|$ does not contain reliable information)? The second key observation is that it can be done using "certifying tournaments." To illustrate this point, consider again the case $k = 2$, where M is a matrix. Split M into two random sub-matrices by sampling the rows w.p. $1/2$. If there is a heavy row, then with probability close to 1, one sub-matrix will have a significantly larger norm then the other. Recall that the method of [32] gives us a $\log n$-approximation. Thus, for very heavy rows, the *ratio* between approximations of norms obtained by the method from [32] will be large. On the other hand, we show that if there are no heavy rows, then such behavior is quite unlikely to be observed many times. Thus, there exists a way to distinguish between the first and the second cases for $(1 - \frac{\epsilon}{\log^2 n})$-significant rows.[2]

The method of certifying tournaments can be generalized to any $s \leq k$ as follows. Let M be a s-dimensional tensor with entries $m_{\mathbf{i}}$ for $\mathbf{i} \in [n]^s$. We "split" M into two "sampled" s-dimensional tensors M^0 and M^1 by randomly sampling the first coordinate. That is, M^1 has entries $m_{\mathbf{i}} H(\mathbf{i}_1)$ and M^0 has entries $m_{\mathbf{i}}(\mathbf{1} - H(\mathbf{i}_1))$, where $H : [n] \mapsto \{0, 1\}$ is pairwise independent and uniform. If there exists a β-approximation algorithm for sampled tensors, and there exists an ϵ-approximation algorithm for Suffix-Sum, $|T_1(M^0)|$ and $|T_1(M^0)|$, then we can approximate L_1 norm of significant hyperplanes. Indeed, if there exists a significant hyperplane M_l of M, then the ratio between β-approximations of $|M^0|$ and $|M^1|$ will be large. If this is the case, the approximation of $T(M^{H(l)})$ is also an ϵ-approximation of $|M_l|$.

To summarize, our main technical Theorem 4.3 proves that it is possible to output a number U such that U is either an approximation of some hyperplane or 0. Further, if there exists a $(1 - \frac{\epsilon}{\beta^2})$-significant hyperplane, then with high probability, U is its approximation. We call such an algorithm an α-ThresholdMax algorithm, for $\alpha = O(\frac{\epsilon}{\beta^2})$.

Indirect Sampling. Many streaming algorithms compute statistics on *sampled* streams, which are random

[2] It is worth noting that the idea of "split-and-compare" is not new. Group testing [22] exploits a similar approach. However, the methods from [22] require ϵ-approximation of L_1; in contrast, we use certifying tournaments to improve the approximation.

subsets of D defined by some randomness $\mathcal{H}$. In many cases, a sampled stream directly corresponds to a collection of sampled entries of a frequency vector. In contrast, subsets of D do not correspond directly to entries M_{Ind}. Thus, our algorithms employ *indirect* sampling, where randomness defines sampled entries of M_{Ind} rather then the entries of a data stream D. We define a Prefix-Zero operator.

Definition 1.13. *Let M be a s-dimensional tensor with entries $m_{\mathbf{i}}, \mathbf{i} \in [n]^s$ and let $H_1, \ldots, H_t, t \leq s$ be hash functions $H_j : [n] \mapsto \{0,1\}$. A* Prefix-Zero *tensor $W(M, H_1, \ldots, H_t)$ a is a s-dimensional tensor with entries $m_{\mathbf{i}} \prod_{l=1}^{t} H_l(\mathbf{i}_l)$.*

Our algorithms work with tensors that are defined by compositions of $\mathcal{F}_{Ind}$, Prefix-Zero and Suffix-Sum. We thus extend the definition of implicitly defined tensors.

Definition 1.6. (Revised) *Let $\mathcal{D}$ be a collection of data streams of size m of elements from domain Ω and let $\mathfrak{H}$ be a collection of hash functions from $[n]$ to $\{0,1\}$. Let $\mathcal{F} : \mathcal{D} \times \mathfrak{H}^t \times [n]^s \mapsto R$ be a fixed function, for some $0 \leq t \leq s$. We say that a s-dimensional tensor M with entries $m_{\mathbf{i}} = \mathcal{F}(D, \mathcal{H}, \mathbf{i}), \mathbf{i} \in [n]^s$ is* implicity defined *by $\mathcal{F}$, given $D \in \mathcal{D}$ and $\mathcal{H} \in \mathfrak{H}^t$. We denote an implicitly defined tensor as $\mathcal{F}(D, \mathcal{H})$.*

Example 1.14. *Consider $k = 2$. Then $\mathcal{F}'(D, H) = W(\mathcal{F}_{Ind}(D), H)$ defines a matrix that represents a collection of rows sampled by a hash function $H : [n] \mapsto \{0,1\}$.*

Generalizing the Method of Indyk and Woodruff [33] to Work on Implicit Vectors. The ThresholdMax algorithm solves the problem that resembles the well-known problem of finding an element with maximal frequency, see, e.g., [17] and [21]. The celebrated method of Indyk and Woodruff [33] uses maximal entries to estimate L_p norms on vectors defined by frequencies. We apply the ideas of [33] to approximate $|AbsoluteVector(M)| = |M|$.

Unfortunately, the method of Indyk and Woodruff [33] is not directly applicable since some basic tools available for frequency vectors (such as L_2 norm approximation) cannot be used. We propose a different algorithm which is still in the same spirit as [33]; it can be found in Section 5. We prove Lemmas 5.5 and 5.3 which state that an existence of an α-ThresholdMax algorithm for an implicitly defined vector V implies

an existence of an (ϵ, δ)-approximation algorithm for $|V|$, with memory increased by an additional factor of $\frac{1}{\alpha} poly(\frac{1}{\epsilon} \log \frac{nm}{\delta})$.

Other Technical Issues. There are several other rather technical issues that need to be resolved. We need to prove that the methods of Indyk [31] and Indyk and McGregor [32] are applicable for k-dimensional tensors that are obtained from M_{Ind} by applying Prefix-Zero and Suffix-Sum operators. The proofs can be found in Section 6. To prove our main theorems, certain properties of the operations on tensors should be established. We prove these in Section 3.

1.5 Related Work

Measuring pairwise independence between two or more random variables is a fundamental problem that touches many areas of computer science. The problems of efficiently testing pairwise, or k-wise, independence were recently considered by Alon, Andoni, Kaufman, Matulef, Rubinfeld and Xie [2]; Alon, Goldreich and Mansour [4]; Batu, Fortnow, Fischer, Kumar, Rubinfeld and White [10]; and Batu, Kumar and Rubinfeld [12]. They addressed the problem of minimizing the number of samples needed to obtain sufficient approximation, when the joint distribution is accessible through a sampling procedure. Unlike the work in [2, 4, 10, 12], in the streaming model, the joint distribution is given by a stream of tuples.

Many exciting results have been reported in the streaming model, including, for example, Alon, Duffield, Lund and Thorup [3]; Alon, Matias and Szegedy [5]; Bagchi, Chaudhary, Eppstein and Goodrich [9]; Bar-Yossef, Jayram, Kumar and Sivakumar [7]; Bar-Yossef, Kumar and Sivakumar [8]; Beame, Jayram and Rudra [13]; Bhuvanagiri, Ganguly, Kesh and Saha [14]; Chakrabarti, Khot and Sun [16]; Charikar, OCallaghan and Panigrahy [18]; Coppersmith and Kumar [19]; Cormode, Datar, Indyk and Muthukrishnan [20]; Datar, Immorlica, Indyk, and Mirrokni [23]; Duffield, Lund and Thorup [24]; Feigenbaum, Kannan, McGregor, Suri and Zhang [25]; Gal and Gopalan [27]; Ganguly [28]; Indyk [31]; Indyk and McGregor [32]; Indyk and Woodruff [33]; Mitzenmacher and Vadhan [37]; Sun and Woodruff [42]; and Szegedy [43]. For a detailed discussion of the streaming model, we refer readers to the excellent surveys of Aggarwal (ed.) [1]; Babcock, Babu, Datar, Motwani and Widom [6]; and Muthukrishnan [38].

In our recent work, [15], we also address the problem of k-wise independence for data stream. In contrast to the current paper, in [15] we study the L_2 norm and use entirely different techniques.

1.6 Roadmap

Section 2 describes the main theorems of the paper. In Section 3 we show some useful properties of Suffix-Sum and Prefix-Zero. Section 4 contains proof of the Tournament algorithm. Section 5 contains a generalization of the ideas of Indyk and Woodruff [33] to implicit vectors. Finally, Section 6 generalizes methods of Indyk [31] and Indyk and McGregor [32] to work with sampled portions of M_{Ind}.

2 Main Theorems

The proof of our result is based on three main steps which are summarized by the following theorems. The remainder of this paper is devoted to establishing these theorems.

Theorem 2.1. ***Dimension Reduction for Implicit Tensors***

Let $s \geq 1$ and let M be a s-dimensional tensor with $poly(n,m)$-bounded entries that is defined by a function $\mathcal{F}$, i.e., $M = \mathcal{F}(D, \mathcal{H})$ where D is a data stream and $\mathcal{H}$ is a fixed randomness. Let $H : [n] \mapsto \{0,1\}$ be an arbitrary fixed hash function. Assume that

1. *There exists an algorithm $\mathfrak{A}(D, \mathcal{H}, H, \delta)$ that, given D and an access to $\mathcal{H}$ and H, in one pass obtains $(\log^k(n), \delta)$-approximation of $|W(M,H)|$;*
2. *There exists an algorithm $\mathfrak{B}(D, \mathcal{H}, H, \epsilon, \delta)$ that, given D and an access to $\mathcal{H}$ and H, in one pass obtains an (ϵ, δ)-approximation of $|T_1(W(M,H))|$;*
3. *Both algorithm require memory $\nu(n,m,\epsilon,\delta) \leq O(\left(\frac{1}{\epsilon}\log\frac{nm}{\delta}\right)^{(30+k)^s})$, beyond the memory required for H and $\mathcal{H}$.*

Then there exists an algorithm that in one pass obtains an (ϵ,δ)-approximation of $|M|$ using memory $\left(\frac{1}{\epsilon}\log\frac{nm}{\delta}\right)^{(30+k)^{s+1}}$.

Proof. Follows from Theorem 4.3, Lemma 5.5, Lemma 5.3 and elementary computations.

Indeed, the assumptions of Theorem 2 imply, by Theorem 4.3, an existence of a $\frac{\epsilon}{\log^{2k}(n)}$-ThresholdMax algorithm (see Definition 4.2) for restricted function $\mathcal{F}' = AbsoluteVector(\mathcal{F}(D,\mathcal{H}))$. The existence of a ThresholdMax algorithm implies, by Lemma 5.3, the existence of a Cover algorithm (see Definition

5.2) for $AbsoluteVector(\mathcal{F}(D,\mathcal{H}))$. The assumption that the entries of M are polynomially bounded and Fact 3.7 imply that the entries of $AbsoluteVector(\mathcal{F}(D,\mathcal{H}))$ are polynomially bounded as well. Thus, by Lemma 5.5, there exists an (ϵ,δ)-approximation algorithm for $|AbsoluteVector(\mathcal{F}(D,\mathcal{H}))|$. Finally, $|AbsoluteVector(\mathcal{F}(D,\mathcal{H}))| = \sum_{i\in[n]} |Hyperplane(\mathcal{F}(D,\mathcal{H}),l)| = |M|$.

After substituting the parameters, the memory required is less than (for sufficiently large n)

$$\frac{1}{\epsilon^{30}}\log(\frac{1}{\delta})\log^{2k+20}(nm)\nu(n,m,\frac{\epsilon^7}{\log^4(nm)},\frac{\epsilon^{17}}{\log^{2k}(n)\log^8(mn)}) \leq \left(\frac{1}{\epsilon}\log\frac{nm}{\delta}\right)^{(30+k)^{s+1}}.$$

□

Theorem 2.2. ***Approximation Theorem for Tensors***

Let M be a k-dimensional tensor with entries bounded by $poly(n,m)$ and implicitly defined by a function $\mathcal{F}(D)$. Assume that

1. *There exists an algorithm $\mathfrak{B}_s(D,H_1,\ldots,H_s)$ (for some $s<k$) that, given D and an access to fixed hash functions $H_1,\ldots,H_s$, in one pass obtains an (ϵ,δ)-approximation of $|T_s(W(M,H_1,\ldots,H_s))|$;*
2. *There exist algorithms $\mathfrak{A}_{s_1,s_2}(D,H_1,\ldots,H_{s_1})$ (for any $0\leq s_2\leq s_1\leq s$) that, given D and an access to H_is, in one pass obtain a $(\log^k(n),\delta)$-approximation of $|T_{s_2}(W(M,H_1,\ldots,H_{s_1}))|$;*
3. *All algorithms use memory bounded by $O(\left(\frac{1}{\epsilon}\log\frac{nm}{\delta}\right)^{20})$, beyond the memory required for H_is.*

Then there exists an algorithm that in one pass obtains an (ϵ,δ)-approximation of $|M|$ using memory $O(\left(\frac{1}{\epsilon}\log\frac{nm}{\delta}\right)^{(30+k)^k})$.

Proof. Define $g(x) = \left(\frac{1}{\epsilon}\log\frac{nm}{\delta}\right)^{(30+k)^{k-x}}$ First, we show that for any $s_1\leq s$ there exists an algorithm $\mathfrak{B}_{s_1}(D,H_1,\ldots,H_{s_1})$ that gives an (ϵ,δ)-approximation of $|T_{s_1}(W(M,H_1,\ldots,H_{s_1}))|$ and uses memory at most $g(s_1)$.

We prove this fact by induction on s_1. For $s_1=s$, the fact follows from the first assumption of Theorem 2.2 since $g(s)\geq\left(\frac{1}{\epsilon}\log\frac{nm}{\delta}\right)^{20}$. For $s_1<s$, denote $\mathcal{F}'(D,H_1,\ldots,H_{s_1}) = T_{s_1}(W(\mathcal{F}(D),H_1,\ldots,H_{s_1})$.

Denote $M' = \mathcal{F}'(D, H_1, \ldots, H_{s_1})$ and let H be an arbitrary hash function. By Corollary 3.4,

$$W(M', H) = W(\mathcal{F}'(D, H_1, \ldots, H_{s_1}), H) = \tag{1}$$

$$W(T_{s_1}(W(\mathcal{F}(D), H_1, \ldots, H_{s_1}), H) = T_{s_1}(W(M, H_1, \ldots, H_{s_1}, H)).$$

Thus, and by the second assumption of the theorem, there exists an algorithm $\mathfrak{A}_{s_1,s_1+1}$ that in one pass obtains a $(\log^k(n), \delta)$-approximation of $|W(M', H)|$ using memory less than or equal to $g(s_1 + 1)$.

Also, by Corollary 3.5 and by (1):

$$T_1(W(M', H)) = T_1(T_{s_1}(W(M, H_1, \ldots, H_{s_1}, H))) = T_{s_1+1}(W(M, H_1, \ldots, H_{s_1}, H)). \tag{2}$$

By induction, there exists an algorithm that gives an (ϵ, δ)-approximation of $|T_{s_1+1}(W(M, H_1, \ldots, H_{s'}, H))| = |T_1(W(M', H))|$ using memory $g(s_1 + 1)$.

M' is implicitly defined by a fixed function $\mathcal{F}'(D, H_1, \ldots, H_s)$. By Fact 3.7, its entries are polynomially bounded. Thus, by (1) and (2), all assumptions of Theorem 2.1 are satisfied for M'. Therefore, there exists an algorithm that gives an ϵ-approximation of $|M'| = |T_{s_1}(W(M, H_1, \ldots, H_{s_1}))|$ using memory $g(s_1)$.

In particular, there exists an algorithm that for any H gives an ϵ-approximation of $|T_1(W(M, H))|$ using $g(1)$. Also, by the second assumption of the theorem, there exists an algorithm that gives a $\log^k(n)$-approximation of $|T_0(W(M, H))| = |W(M, H)|$. Thus, we can apply Theorem 2.1 for M and obtain an ϵ-approximation of $|M|$. The resulting memory usage will be $O((\frac{1}{\epsilon}\log\frac{nm}{\delta})^{(30+k)^k})$. □

The following lemmas are proven in Section 6.

Lemma 2.3. *There exists an algorithm $\mathfrak{B}_{k-1}$ that, given a data stream D and an access to hash functions $H_1, \ldots, H_{k-1}$, in one pass obtains an ϵ-approximation of $|T_{k-1}(W(M_{Ind}, H_1, \ldots, H_{k-1}))|$ using memory $O(\frac{1}{\epsilon^2}\log\frac{1}{\delta}\log\frac{nm}{\epsilon\delta})$.*

Lemma 2.4. *There exists an algorithm $\mathfrak{A}_{s_1,s_2}$ (for any $0 \le s_2 \le s_1 \le k$) that, given a data stream D and an access to hash functions $H_1, \ldots, H_{s_1}$, in one pass obtains a $\log^k n$-approximation of $|T_{s_2}(W(M_{Ind}, H_1, \ldots, H_{s_1}))|$ using memory $O(\log(nm)\log\frac{1}{\delta})$.*

Theorem 2.5. ***Main Theorem*** *Let $k \geq 2$ be a constant, and let D be a stream of k-tuples from $[n]^k$. For any $0 < \epsilon < 1$, there exists an algorithm that makes a single pass over D and returns an (ϵ, δ)-approximation of the statistical distance between product and joint distribution (see Definition 1.1) using memory $O(\left(\frac{1}{\epsilon}\log(\frac{nm}{\delta})\right)^{(30+k)^k})$.*

Proof. By Lemma 2.3 and Lemma 2.4, the algorithms required by Theorem 2.2 exist for M_{Ind}. Also, by Fact 3.7, the entries of M_{Ind} are polynomially bounded. Thus all assumptions of Theorem 2.2 are true for M_{Ind}. Applying Theorem 2.2 to M_{Ind}, we obtain the main result. □

3 Properties of Tensors

We prove the following useful facts about Suffix-Sum and Prefix-Zero operations.

Fact 3.1. *Let M be a t-dimensional tensor and $0 \leq s \leq t$. Then*

$$W(T_s(M), H) = T_s(W(M, H_1 = \mathbf{1}, \ldots, H_s = \mathbf{1}, H)).$$

Proof. Denote by $m_{\mathbf{w}}$ (for $\mathbf{w} \in [n]^t$) the $\mathbf{w}$-th entry of M. For any $\mathbf{i} \in [n]^{t-s}$, denote by $a_{\mathbf{i}}$ the entry of $T_s(M)$. By Definition 1.12:

$$a_i = \sum_{\mathbf{j} \in [n]^s} m_{(\mathbf{j},\mathbf{i})}.$$

Denote by $b_{\mathbf{i}}$ the entry of $W(T_s(M), H)$. By Definitions 1.12 and 1.13:

$$b_{\mathbf{i}} = H(\mathbf{i}_1)a_{\mathbf{i}} = \sum_{\mathbf{j} \in [n]^s} m_{(\mathbf{j},\mathbf{i})}H(\mathbf{i}_1).$$

Denote by $c_{\mathbf{i}}$ the $\mathbf{i}$-th entry of $T_s(W(M, H_1 = \mathbf{1}, \ldots, H_s = \mathbf{1}, H))$. By Definitions 1.12 and 1.13:

$$c_{\mathbf{i}} = \sum_{\mathbf{j} \in [n]^s} m_{(\mathbf{j},\mathbf{i})}H(\mathbf{i}_1).$$

Thus, for any $\mathbf{i}$, $b_{\mathbf{i}} = c_{\mathbf{i}}$ and the fact is correct. □

Fact 3.2. *Let M be a t-dimensional tensor and let $0 \leq s < t$. Then $T_1(T_s(M)) = T_{s+1}(M)$.*

Proof. Denote by $m_{\mathbf{w}}$ (for $\mathbf{w} \in [n]^t$) the $\mathbf{w}$-th entry of M. For $\mathbf{j} \in [n]^{t-s}$ denote $b_{\mathbf{j}}$ to be an entry of $T_s(M)$. By Definition 1.12:

$$b_{\mathbf{j}} = \sum_{\mathbf{u} \in [n]^s} m_{(\mathbf{u},\mathbf{j})}.$$

For every $\mathbf{i} \in [n]^{t-s-1}$, denote by $c_{\mathbf{i}}$ the entry of $T_1(T_s(M))$. By Definition 1.12:

$$c_{\mathbf{i}} = \sum_{l \in [n]} b_{(l,\mathbf{i})} = \sum_{l \in [n]} \sum_{\mathbf{u} \in [n]^s} m_{(\mathbf{u},(l,\mathbf{i}))} = \sum_{l \in [n]} \sum_{\mathbf{u} \in [n]^s} m_{((\mathbf{u},l),\mathbf{i})} = \sum_{\mathbf{v} \in [n]^{s+1}} m_{(\mathbf{v},\mathbf{i})}.$$

For any $\mathbf{i} \in [n]^{t-s-1}$ denote by $a_{\mathbf{i}}$ the entry of $T_{s+1}(M)$. By Definition 1.12:

$$a_{\mathbf{i}} = \sum_{\mathbf{v} \in [n]^{s+1}} m_{(\mathbf{v},\mathbf{i})}.$$

Thus, for any $\mathbf{i}$, $a_{\mathbf{i}} = c_{\mathbf{i}}$ and the fact is correct. □

Fact 3.3. *Let M be a t-dimensional tensor, let $s \leq t$ and let $H_1, \ldots, H_s$ and $G_1, \ldots, G_s$ be hush functions. Then*

$$W(M, H_1G_1, \ldots, H_sG_s) = W(W(M, H_1, \ldots, H_s), G_1, \ldots, G_s))$$

Corollary 3.4. *Let M be a t-dimensional tensor and let $0 \leq s < t$. Let $M' = T_s(W(M, H_1, \ldots, H_s))$. Then*

$$W(M', H) = T_s(W(M, H_1, \ldots, H_s, H)).$$

Proof. Denote $M'' = W(M, H_1, \ldots, H_s)$. Then by Fact 3.1:

$$W(M', H) = W(T_s(M''), H) = T_s(W(M'', G_1 = \mathbf{1}, \ldots, G_k = \mathbf{1}, H)).$$

Also by Fact 3.3:

$$W(M'', G_1, \ldots, G_s, H) = W(W(M, H_1, \ldots, H_s, \mathbf{1}), G_1, \ldots, G_s, H) = W(M, H_1, \ldots, H_s, H).$$

□

Corollary 3.5. *Let M be a t-dimensional tensor and let $0 \leq s < t$. Let $M' = T_s(W(M, H_1, \ldots, H_s))$. Then*

$$T_1(M', H)) = T_{s+1}(W(M, H_1, \ldots, H_s, H)).$$

Proof. By Fact 3.2 and Corollary 3.4:

$$T_{s+1}(W(M, H_1, \ldots, H_s, H)) = T_1(T_s(W(M, H_1, \ldots, H_s, H))) = T_1(W(M', H)).$$

□

Fact 3.6. *Let M be an arbitrary s-dimensional tensor, let M_l be $(1-\epsilon/2)$-significant hyperplane of M, $M_l = Hyperplane(M, l)$, and let $M' = T_1(M)$. Then $|M'|$ is an ϵ-approximation of $|M_l|$.*

Proof. We have

$$|M'| = \sum_{\mathbf{i}\in[n]^{s-1}} |\sum_{j\in[n]} m_{(j,\mathbf{i})}| \leq \sum_{\mathbf{i}\in[n]^{s-1}} \sum_{j\in[n]} |m_{(j,\mathbf{i})}| = |M| \leq \frac{1}{1-\epsilon/2}|M_l| \leq (1+\epsilon)|M_l|.$$

On the other hand,

$$|M'| = \sum_{\mathbf{i}\in[n]^{s-1}} |\sum_{j\in[n]} m_{(j,\mathbf{i})}| \geq \sum_{\mathbf{i}\in[n]^{s-1}} (|m_{(l,\mathbf{i})}| - \sum_{j\in[n],j\neq l} |m_{(j,\mathbf{i})}|) = |M_l| - (|M| - |M_l|) =$$

$$\geq (2 - \frac{1}{1-\epsilon/2})|M_l| \geq (1-\epsilon)|M_l|.$$

□

Fact 3.7. *The following is correct:*

1. *Let M be a s-dimensional tensor with polynomially bounded (in n and m) entries for $s \leq k$. Let M' be a tensor obtained from M by an arbitrary composition of Prefix-Zero, AbsoluteVector, Hyperplane and Suffix-Sum operators. Then the entries of M' are polynomially bounded.*

2. *All entries of M_{Ind} are integers with absolute values bounded by $2m^k$ and thus claim 1 is true for M_{Ind}.*

Proof. The first claim follows from the fact that the entries of M' are sums of disjoint subsets of M and that the number of entries in M is bounded by n^k. The second claim follows from Definition 1.7. □

4 Certifying Tournaments

Algorithm 4.1. *TensorTournament($D, \mathcal{H}, H, \epsilon$)*

1. *Repeat in parallel* $O(\frac{\log \frac{1}{\delta}}{p})$ *times where* $p = 1 - \sqrt{1 - \epsilon/2}$.
   (a) *Generate 2-wise independent random hash function* Z *from* $[n]$ *to* $\{0, 1\}$ *such that* $Z(i) = 0$ *w.p.* 0.5. *Denote* $Z_1 = HZ, \;\; Z_0 = H(1 - Z)$.
   (b) *Compute in a single pass over* D *for* $i = 0, 1$: $\mathfrak{t}_i = \mathfrak{A}(D, \mathcal{H}, Z_i, \epsilon, \delta')$, *where* $\delta' = \frac{p\epsilon}{4 \log \frac{1}{\delta}}$.
   (c) *Simultaneously (in the same pass), compute* $\mathfrak{l}_i = \mathfrak{B}(D, \mathcal{H}, Z_i, \delta')$.
   (d) *Put* $\mathfrak{u}_i = \max\{\frac{\mathfrak{l}_i}{\beta}, \mathfrak{t}_i, 0\}, \;\; i = 0, 1$.
   (e) *Define* $\lambda' = (1 + \epsilon)\lambda$, *where* λ *is the constant from Lemma 4.4,* $\lambda = (1 + \frac{2(1-\epsilon)^{1/4}}{1-(1-\epsilon)^{1/4}})$.
   (f) *Compute*
   $$U' = \begin{cases} \mathfrak{u}_1, & \text{if } \mathfrak{u}_1 \geq \lambda' \beta^2 \mathfrak{u}_0, \\ \mathfrak{u}_0, & \text{if } \mathfrak{u}_0 \geq \lambda' \beta^2 \mathfrak{u}_1, \\ 0, & \text{otherwise}. \end{cases}$$
2. *Output* U *to be the minimum of all* U'*s.*

Definition 4.2. *Let* $\mathcal{F}$ *be a fixed function that defines implicit vectors, given a data stream and a fixed randomness and denote* $V = \mathcal{F}(D, \mathcal{H})$ *as a vector with entries* v_i. *For* $\alpha > 0.5$, *an* α-ThresholdMax algorithm for restricted $\mathcal{F}$ *is an algorithm that receives as an input a data stream* D *and an access to a randomness* $\mathcal{H}$ *and a random function* $H : [n] \mapsto \{0, 1\}$, *and in one pass over* D *returns* $U \geq 0$ *such that w.p. at least* $1 - \delta$*:*

1. *If* $U > 0$ *then* U *is an* ϵ*-approximation of* $|v_i|$ *for some* i *with* $H(i) = 1$.
2. *If*[3] $|VH| > 0$ *and there exists* i *such that* $H(i) = 1$ *and* $|v_i| \geq (1 - \alpha)|VH|$ *then* U *is an* ϵ*-approximation of* $|v_i|$.

Theorem 4.3. *Let* H *be a fixed hash function defined as above and let* $\epsilon \leq 0.1$. *Let* M *be a* s*-dimensional tensor implicitly defined by a fixed function* $\mathcal{F}$, *stream* D *and randomness* $\mathcal{H}$, $M = \mathcal{F}(D, \mathcal{H})$. *If there exist:*

- *An algorithm* $\mathfrak{A}(D, \mathcal{H}, H, \delta)$ *that in one pass obtains* (β, δ)*-approximation of* $|W(M, H)|$ *using memory* $\mu_1(n, m, \epsilon, \delta)$*;*

[3] Here and thenceforth we denote by VH a vector with entries $v_i H(i)$, $i \in [n]$

- *An algorithm $\mathfrak{B}(D, \mathcal{H}, H, \epsilon, \delta)$ that in one pass over D obtains an (ϵ, δ)-approximation of $|T_1(W(M, H))|$ using memory $\mu_2(n, m, \epsilon, \delta)$;*

Let $\alpha = \frac{\epsilon}{64\beta^2}$. Then the Algorithm $TensorTournament(D, \mathcal{H}, H, \epsilon)$ is an α-ThresholdMax algorithm for restricted $\mathcal{F}'$ (see Definition 4.2), where $\mathcal{F}'(D, \mathcal{H}) = AbsoluteVector(\mathcal{F}(D, \mathcal{H}))$. The algorithm makes a single pass over D and uses memory

$$O(\frac{1}{\epsilon}\log\frac{1}{\delta}(\mu_1(n, m, \epsilon/3, \delta\epsilon/\log(1/\delta)) + \mu_2(n, m, \epsilon/3, \delta\epsilon/\log(1/\delta)) + \log nm).$$

Proof.

Denote $M^t = W(M, Z_t)$ for $t = 0, 1$. Let $M_i = Hyperplane(M, i)$ for $i \in [n]$ and let V' to be a vector with elements $|M_i|$. By Definition 1.11, $V' = \mathcal{F}'(D, \mathcal{H})$. Further, let V be a vector with entries $v_i = |M_i|H(i)$. We prove that the algorithm satisfies two conditions of Definition 4.2 for the ThresholdMax algorithm for V and H.

Proof of the first condition of Definition 4.2

We prove the following stronger statements which imply the first condition of Definition 4.2:

I. If there is no $(1-\epsilon)$-significant entry v_l then, w.p. at least $1 - \frac{\delta}{3}$, $U = 0$.

II. If $|V| > 0$ and there is a $(1-\epsilon)$-significant entry v_l then, w.p. at least $1 - \frac{\delta}{3}$, either $U = 0$ or U is a 3ϵ-approximation of $|v_l|$.

Proof of statement I

By definitions of $\mathfrak{B}, \mathfrak{A}$, we have w.p. at least $1 - 8\delta'$ for $t = 0, 1$: $\mathfrak{u}_t \geq \frac{\mathfrak{l}_t}{\beta} \geq \frac{|M^t|}{\beta^2}$; and $\mathfrak{t}_t \leq (1+\epsilon)|T_1(M^t, H)| \leq (1+\epsilon)|M^t|$; and $\frac{\mathfrak{l}_t}{\beta} \leq |M^t|$. Thus,

$$\frac{|M^t|}{\beta^2} \leq \mathfrak{u}_t \leq (1+\epsilon)|M^t|. \tag{3}$$

Following the terminology of Lemma 4.4, we define $X = |M^1|$ and $Y = |M^0|$. We have the following relations:

$$|V| = \sum_i v_i = \sum_i H(i)|M_i| = \sum_{i\in[n]} H(i) \sum_{\mathbf{j}'\in[n]^{s-1}} |m_{(i,\mathbf{j}')}| = |W(M, H)|,$$

$$X = |M^1| = \sum_{\mathbf{j}\in[n]^s} Z(\mathbf{j}_1)H(\mathbf{j}_1)|m_{\mathbf{j}}| = \sum_{i\in[n]} Z(i)H(i) \sum_{\mathbf{j}'\in[n]^{s-1}} |m_{(i,\mathbf{j}')}| = \sum_i Z(i)H(i)|M_i| = \sum_i Z(i)v_i,$$

and similarly

$$Y = |M^0| = \sum_i (1 - Z(i))H(i)|M_i| = |V| - X = |V| - |M^1|. \tag{4}$$

By statement **I**, for all i, $v_i < (1-\epsilon)|V|$. Thus we can apply Lemma 4.4. We have:

$$P((|M^0| \geq \lambda|M^1|) \cup (|M^1| \geq \lambda|M^0|)) = P((X \geq \lambda Y) \cup (Y \geq \lambda X)) \leq \sqrt{1-\epsilon}.$$

Let Υ be the event $(\mathfrak{u}_0 \geq \lambda'\beta^2\mathfrak{u}_1) \cup (\mathfrak{u}_1 \geq \lambda'\beta^2\mathfrak{u}_0)$. Let Φ be the event that $\frac{|M^t|}{\beta^2} \leq \mathfrak{u}_t \leq (1+\epsilon)|M^t|$ for both values of t. We have $P(\Upsilon) \leq P(\Upsilon, \Phi) + P(\bar{\Phi})$. By (3), we have $P(\bar{\Phi}) \leq 8\delta'$. Also, events $\mathfrak{u}_0 \geq \lambda'\beta^2\mathfrak{u}_1$ and Φ imply that $|M^0| \geq \lambda|M^1|$; indeed:

$$|M^0| \geq \frac{\mathfrak{u}_0}{(1+\epsilon)} \geq \frac{\lambda'}{1+\epsilon}\beta^2\mathfrak{u}_1 \geq \lambda|M^1|.$$

Thus we have

$$P(\Upsilon, \Phi) \leq P((|M^0| \geq \lambda|M^1|) \cup (|M^1| \geq \lambda|M^0|)) \leq \sqrt{1-\epsilon}.$$

We summarize that if no $(1-\epsilon)$-significant v_i exists, then

$$P(U' \neq 0) \leq P(\Upsilon) \leq \sqrt{1-\epsilon} + O(\delta') \leq \sqrt{1-\epsilon/2}.$$

Recall that the number of repetitions is $O(\frac{1}{p}\log 1/\delta)$, where $p = 1 - \sqrt{1-\epsilon/2}$. Thus $P(U \neq 0) \leq (1-p)^{\frac{1}{p}\log\frac{3}{\delta}} \leq \frac{\delta}{3}$.

Proof of statement II

Let v_l be a $(1-\epsilon)$-significant entry of V. Assume, w.l.o.g., that for one execution of the main cycle of the $Tournament$ algorithm, $Z(l) = 0$. Statement **II** implies $|V| > 0$ which implies $v_l = |M_l|H(l) > 0$ which implies $(1 - Z(l))H(l) = 1$. Thus, $|Hyperplane(M^0, l)| = |Hyperplane(W(M, (\mathbf{1} - Z)H), l)| = |M_l| = v_l$. Therefore by (4), $|Hyperplane(M^0, l)| = v_l \geq (1-\epsilon)|V| \geq (1-\epsilon)|M^0|$, i.e., the l-th hyperplane of M^0 is $(1-\epsilon)$-significant. By Fact 3.6, $|T(M^0)|$ is an 2ϵ-approximation of $|M_l|$. By the assumptions of the theorem, $\mathfrak{B}$ returns an ϵ-approximation of $|T(M^0)|$. Thus, $\mathfrak{t}_0$ is a 3ϵ-approximation of $|M_l|$, w.p. at least $1 - \delta'$, in which case

$$\mathfrak{u}_0 \geq \mathfrak{t}_0 \geq (1 - 3\epsilon)|M_l|.$$

Also, by the assumption of Theorem 4.3, w.p. at least $1-\delta'$, we have $\frac{\mathfrak{l}_0}{\beta} \leq |M^0|$. Thus

$$\mathfrak{u}_0 = \max\{\frac{\mathfrak{l}_0}{\beta}, \mathfrak{t}_0, 0\} \leq \max\{|M^0|, (1+3\epsilon)|M_l|\} \leq (1+3\epsilon)|M_l|.$$

On the other hand, w.p. at least $1-2\delta'$

$$\mathfrak{u}_1 = \max\{\frac{\mathfrak{l}_1}{\beta}, \mathfrak{t}_1, 0\} \leq \max\{|M^1|, (1+\epsilon)|M^1|\} = (1+\epsilon)|M^1|.$$

But since $Z_s(l) = 0$ we have by (4):

$$|M^1| = |V| - |M^0| \leq |V| - |Hyperplane(M^0, l)| = |V| - v_l \leq \frac{\epsilon}{1-\epsilon}|M_l|.$$

Combining all of the above computations, we conclude that w.p. at least $1-4\delta'$ (for sufficiently small ϵ, e.g., $\epsilon \leq 0.1$):

$$\mathfrak{u}_1 \leq (1+\epsilon)|M_1| \leq \frac{\epsilon(1+\epsilon)}{1-\epsilon}|M_l| \leq \frac{\epsilon(1+\epsilon)}{(1-\epsilon)(1-3\epsilon)}\mathfrak{u}_0 < \lambda'\mathfrak{u}_0.$$

Thus, U' is equal to either 0 or $\mathfrak{u}_0$ w.p. at least $1-4\delta'$. Recall simultaneously $\mathfrak{u}_0$ is a 3ϵ-approximation of $|M_l| = v_l$. The same inequality is true if $Z(l) = 1$. By union bound, w.p. at least $1-\Omega(\frac{\log\frac{1}{\delta}}{p}\delta') = 1-\Omega(\delta)$, U is either 0 or a 3ϵ-approximation of v_l.

Proof of the second condition of Definition 4.2

Finally, consider the case when v_l is a $(1-\alpha)$-significant entry of V. Consider the case when $Z(l) = 0$. Repeating the arguments from the proof of statement **II**, we have, w.p. at least $1-4\delta'$, $\mathfrak{u}_0$ is a 3ϵ-approximation of v_l and

$$\mathfrak{u}_1 \leq (1+\epsilon)|M^1| \leq (1+\epsilon)\frac{\alpha}{(1-\alpha)}v_l \leq 4\alpha v_l.$$

Therefore,

$$\mathfrak{u}_0 \geq (1-3\epsilon)v_l \geq \frac{(1-3\epsilon)}{4\alpha}\mathfrak{u}_1 \geq \lambda'\beta^2\mathfrak{u}_1.$$

Thus, w.p. $1-4\delta'$, $U' = \mathfrak{u}_0 = (1\pm 3\epsilon)v_l$. The same is true when $Z(l) = 1$. Thus, U is a 3ϵ-approximation of v_l w.p. at least $1-\Omega(\delta)$.

Conclusion and memory analysis

Since both conditions of Definition 4.2 are met (substituting ϵ with $\epsilon/3$), we conclude that $TensorTournament$ is an α-ThresholdMax algorithm for restricted $\mathcal{F}'$. Let us count the memory needed for a single iteration of the main cycle of the algorithm. To generate pairwise independent Z, we need $O(\log n)$ bits. In addition, we need $\mu_1 + \mu_2$ for the algorithms $\mathfrak{B}$ and $\mathfrak{A}$ and $O(\log nm)$ bits to keep the auxiliary variables. Thus, in total we need memory

$$O(\frac{1}{\epsilon}\log\frac{1}{\delta}(\mu_1(n,m,\epsilon/3,\delta\epsilon/\log(1/\delta)) + \mu_2(n,m,\epsilon/3,\delta\epsilon/\log(1/\delta)) + \log nm).$$

Recall that we do not count memory required to store $\mathcal{H}$ and H. □

Lemma 4.4. *Let V be a n-dimensional vector with non-negative entries $v_i \geq 0, i \in [n]$. Let Z be 2-wise independent random hash functions from $[n]$ to $\{0,1\}$, such that $P(Z(i)=1) = 0.5$. Let $X = \sum_i v_i Z(i)$, and $Y = L_1(V) - X$. If there exists $\epsilon > 0$ such that for all i $v_i < (1-\epsilon)L_1(V)$, then for $\lambda = \lambda(\epsilon) \geq 1 + \frac{2(1-\epsilon)^{1/4}}{1-(1-\epsilon)^{1/4}}$ we have*

$$P((X \geq \lambda Y) \cup (Y \geq \lambda X)) \leq \sqrt{1-\epsilon}.$$

Proof. Clearly, $E(X) = L_1(V)/2$. Further, by 2-wise independency of Z, we have

$$E(X^2) = E((\sum_i v_i Z(i))^2) = \frac{1}{2}\sum_i v_i^2 + \frac{1}{4}\sum_{i\neq j} v_i v_j = \frac{1}{4}\sum_i v_i^2 + E(X)^2.$$

Thus, by the assumption that $v_i < (1-\epsilon)L_1(V)$, we have:

$$Var(X) = E(X^2) - E(X)^2 = \frac{1}{4}\sum_i v_i^2 \leq \frac{1-\epsilon}{4}L_1(V)^2.$$

Thus, $\sigma(X) \leq \frac{\sqrt{1-\epsilon}}{2}L_1(V)$. Note that event $X \geq \lambda Y$ is equivalent to the event $X - E(X) \geq \frac{\lambda-1}{2(\lambda+1)}L_1(V)$ and event $Y \geq \lambda X$ is equivalent to the event $E(X) - X \geq \frac{\lambda-1}{2(\lambda+1)}L_1(V)$. Thus

$$P((X \geq \lambda Y) \cup (Y \geq \lambda X)) = P(|E(X) - X| \geq \frac{\lambda-1}{2(\lambda+1)}L_1(V)) \leq$$

$$P(|E(X) - X| \geq \frac{\lambda-1}{2(\lambda+1)}\frac{2}{\sqrt{1-\epsilon}}\sigma(Y)) \leq$$

$$(1-\epsilon)\left(\frac{\lambda+1}{\lambda-1}\right)^2 \leq \sqrt{1-\epsilon}.$$

for $\lambda \geq 1 + \frac{2(1-\epsilon)^{1/4}}{1-(1-\epsilon)^{1/4}}$.

Note that if there is at most one strictly positive v_i, then $P((X \geq \lambda Y) \cup (Y \geq \lambda X)) = 1$ seems to contradict our lemma. However, in this case, there exists v_i, such that $v_i = L_1(V)$, and thus the assumption of the lemma does not hold. Generally, the assumptions imply that there exists at least $\frac{1}{1-\epsilon}$ strictly positive entries v_i. □

5 Approximating L_1 Norms of Implicit Vectors

Definition 5.1. *Let V with $v_i \geq 0$ be a vector from R^n. A set $\mathcal{U}$ of positive numbers is an* ϵ-cover *of V if:*

1. *All elements of $\mathcal{U}$ are ϵ-approximations of distinct and positive coordinates from V. I.e., there is a one-to-one mapping ρ from the set $\mathcal{U}$ to a subset $S' \subseteq [n]$ such that for all $U \in \mathcal{U}$, U is an ϵ-approximation of $v_{\rho(U)}$.*

2. *$\mathcal{U}$ contains ϵ-approximations of all ϵ-significant elements of V. I.e., for all v_i such that $v_i \geq \epsilon|V|$, it is true that $i \in S'$.*

The size *of the cover is $|\mathcal{U}|$.*

Definition 5.2. *Let $\mathcal{F}$ be a fixed function that implicitly defines vectors, given a data stream D and a fixed randomness $\mathcal{H}$. Denote $V = \mathcal{F}(D, \mathcal{H})$. A* Cover algorithm for restricted $\mathcal{F}$ *is an algorithm that receives as an input a data stream D, an access to a randomness $\mathcal{H}$ and a random function $H : [n] \mapsto \{0, 1\}$ and an ϵ and δ. The algorithm makes a single pass over D and w.p. at least $1 - \delta$, returns an ϵ-cover of vector with entries $v_i H(i)$.*

5.1 Witnessing ϵ-Significant Hyperplanes

Lemma 5.3. *Let $\mathcal{F}$ be a fixed function that implicitly defines vectors, given a data stream D and a fixed randomness $\mathcal{H}$. An existence of α-ThresholdMax algorithm for restricted $\mathcal{F}$ that uses memory $\mu(n, m, \epsilon, \delta)$ implies an existence of a Cover algorithm for restricted $\mathcal{F}$ for any ϵ. The Cover algorithm uses memory $O(\frac{1}{\epsilon^2\delta\alpha}(\mu(n, m, \epsilon, \delta^2\epsilon^2\alpha) + \log nm))$.*

Proof. Denote by $\mathfrak{L}_\alpha(D, \mathcal{H}, H, \epsilon, \delta)$ the existing α-ThresholdMax algorithm for restricted $\mathcal{F}$.

Using $\mathfrak{L}_\alpha$ we construct the following algorithm. Let $\epsilon' = \epsilon^2\delta/3$ and $\varrho = \lceil \frac{1}{\epsilon'\alpha} \rceil$. Let G be a pairwise independent random hash function from $[n]$ to $[\varrho]$ that is independent of $\mathcal{H}$ and H. For $s \in [\varrho]$, define

function F_s as $F_s(i) = \mathbf{1}_{G(i)=s}$ and execute, in parallel for all s, $\mathfrak{L}_\alpha(D, \mathcal{H}, HF_s, \epsilon, \delta/\varrho)$. Let U_s be the output of s-th ran of $\mathfrak{L}_\alpha$. The output of our new algorithm is a set of all strictly positive U_s. We show below that the output is indeed ϵ-cover of V with probability at least $1 - \delta$.

Let $V = \mathcal{F}(D, \mathcal{H})$ be a vector with entries v_i and let V_s be a vector with entries $v_{s,i} = v(i)F_s(i)$. By the union bounds and by the definition of α-ThresholdMax algorithm, w.p. at least $1 - \delta$, every positive U_s is an ϵ approximation of $|v_{i_s}|$ for some i_s with $H(i_s)F_s(i_s) = 1$. But this implies that U_s is an approximation of $|v_i|$ with $H(v_i) = 1$. Since G splits $[n]$ into disjoint subsets, the output of our algorithm corresponds to ϵ-approximations of absolute values of a set of distinct entries of V. I.e., the first condition of ϵ-cover is correct.

To show that the second condition is true as well, let S_ϵ be set of all is such that $|v_i H(i)| \geq \epsilon |VH| > 0$. Consider a fixed $i \in S_\epsilon$. Let

$$X_i = |VHF_{G(i)}| - |v_i| = \sum_{j \neq i} |v_j| H(j) F_{G(i)}(j) \geq 0.$$

By pairwise independency of G:

$$E(X_i) = \sum_{j \neq i} |v_j| H(j) P(G(j) = G(i)) \leq \frac{|VH|}{\varrho}.$$

Let Ψ_i be the event that $X_i > \frac{\epsilon}{\varrho\epsilon'}|VH|$; by Markov inequality $P(\Psi_l) \leq \frac{\epsilon'}{\epsilon}$. Note that if Ψ_i does not happen, then

$$|VHF_{G(i)}| - |v_i| \leq \frac{\epsilon}{\varrho\epsilon'}|VH| \leq \frac{1}{\varrho\epsilon'}|v_i| \leq \alpha |v_i|,$$

in which case $|v_i| \geq (1 - \alpha)|VHF_{G(i)}|$. Let Γ_l be the event that $U_{G(i)}$ is not an ϵ-approximation of $|v_i|$. By the properties of algorithm $\mathfrak{L}_\alpha$, $P(\Gamma_i | \bar{\Psi}_i) \leq \frac{\delta}{\varrho}$. Thus

$$P(\Gamma_i) \leq P(\Gamma_i | \bar{\Psi}_i) + P(\Psi_i) \leq \frac{\delta}{\varrho} + \frac{\epsilon'}{\epsilon}.$$

Finally, let $\Phi_{i,j}$ be the event where there is a collision between i and j. By pairwise independence of G, $P(\Phi_{i,j}) = \frac{1}{\varrho}$, and thus the probability of collisions for ϵ-significant entries is bounded by $\frac{1}{\epsilon^2\varrho}$. Thus, the probability that the output of the algorithm does not meet the second condition of ϵ-cover is bounded by

$$P((\cup_{i \in S_\epsilon} \Gamma_i) \cup (\cup_{i,j \in S_\epsilon} \Phi_{i,j})) \leq \frac{\delta}{\varrho\epsilon} + \frac{\epsilon'}{\epsilon^2} + \frac{1}{\varrho\epsilon^2} \leq \delta.$$

□

5.2 The ϵ-Approximation

Definition 5.4. *Let $\mathcal{F}$ be a fixed function that defines an implicit vector $V = \mathcal{F}(D, \mathcal{H})$, given D and a randomness $\mathcal{H}$, as in Definition 1.6. An algorithm that receives as an input a data stream D and an access to a randomness $\mathcal{H}$ and in one pass over D returns an (ϵ, δ)-approximation of $|\mathcal{F}(D, \mathcal{H})|$ is called an* (ϵ, δ)-approximation algorithm for $L_1(\mathcal{F})$.

The main goal of this section is to prove

Lemma 5.5. *Let $\mathcal{F}$ be a fixed function that defines an implicit vector $V = \mathcal{F}(D, \mathcal{H})$, given D and a randomness $\mathcal{H}$. Assume that V has non-negative entries bounded by $poly(n, m)$. Then the existence of Cover algorithm $\mathfrak{Q}(D, \mathcal{H}, H, \epsilon, \delta)$ for restricted $\mathcal{F}$ (see Definition 5.2) that uses memory $\mu(n, m, \epsilon, \delta)$ implies an existence of an $(\epsilon, 2/3)$-approximation algorithm for $L_1(\mathcal{F})$ (Definition 5.4) that uses memory*

$$O\left(\frac{1}{\epsilon}\log(n)\mu(n, m, \frac{\epsilon^7}{\log^3(nm)}, \frac{\epsilon}{\log(nm)}) + \frac{1}{\epsilon^2}\log^2(nm)\right).$$

5.2.1 Notations

In this section, let $0 < \epsilon < 1$ be a constant, Define

$$\mathfrak{a} = O(\log_{(1+\epsilon)} n), \quad \mathfrak{b} = O(\log_{(1+\epsilon)} nm), \quad \chi' = 10(\mathfrak{a} + \mathfrak{b}), \quad \chi = \lceil \frac{16}{\epsilon^3}\chi' \rceil,$$

$$Q = \lceil \frac{20\chi'}{\epsilon^2} \rceil, \quad \zeta = (1+\epsilon)^{1/Q} - 1, \quad \mathfrak{c} = \min\{\frac{\zeta}{2(1+\zeta)}, \frac{\epsilon}{4\chi\chi'^2}\}.$$

For $x > \chi$, let $\mathfrak{f}_\chi(x)$ be an integer such that $\chi(1+\epsilon)^{\mathfrak{f}_\chi(x)-1} < x \leq \chi(1+\epsilon)^{\mathfrak{f}_\chi(x)}$ i.e., $\mathfrak{f}_\chi(x) = \lceil \log_{(1+\epsilon)} \frac{x}{\chi} \rceil$. It is easy to see that for $x > \chi$ we have $\mathfrak{f}_\chi(x) > 0$.

5.2.2 Technical Lemmas

Let n_{rest} be such that $n \geq n_{rest} > \chi$. For any $j = [\mathfrak{a}]$ and for every $i \in [n_{rest}]$, let $X_{i,j}$ be pairwise independent zero-one random variables with $P(X_{i,j} = 1) = \frac{1}{(1+\epsilon)^j}$. Let $Y_j = \sum_{i \in [n_{rest}]} X_{i,j}$.

Fact 5.6.

$$P(\frac{\chi}{(1+\epsilon)^2} < Y_{\mathfrak{f}_\chi(n_{rest})} \leq (1+3\epsilon)\chi) \geq 1 - \frac{1}{\chi'}.$$

Proof. Let $j_0 = \mathfrak{f}_\chi(n_{rest})$; note that $j_0 > 0$ since $n_{rest} > \chi$. We have $E(Y_{j_0}) = \frac{n_{rest}}{(1+\epsilon)^{j_0}}$. and, by pairwise independency of $X_{j_0,i}$:

$$Var(Y_{j_0}) = n_{rest}Var(X_{j_0,1}) = \frac{n_{rest}}{(1+\epsilon)^{j_0}}(1 - \frac{1}{(1+\epsilon)^{j_0}}) \leq \frac{n_{rest}}{(1+\epsilon)^{j_0}} \leq \chi.$$

Let $\epsilon' = \frac{\epsilon}{2}$; we have, by Chebyshev's inequality:

$$P\left(|Y_{j_0} - \frac{n_{rest}}{(1+\epsilon)^{j_0}}| \geq \epsilon' \frac{n_{rest}}{(1+\epsilon)^{j_0}}\right) \leq Var(Y_{j_0}) \left(\frac{(1+\epsilon)^{j_0}}{\epsilon' n_{rest}}\right)^2 \leq$$

$$\frac{1}{\epsilon'^2}\chi(\frac{(1+\epsilon)^{j_0}}{n_{rest}})^2 \leq \frac{1}{\epsilon'^2}\frac{(1+\epsilon)^2}{\chi} \leq \frac{1}{\chi'}.$$

Also, $|Y_{j_0} - \frac{n_{rest}}{(1+\epsilon)^{j_0}}| < \epsilon' \frac{n_{rest}}{(1+\epsilon)^{j_0}}$ implies

$$Y_{j_0} < (1+\epsilon')\frac{n_{rest}}{(1+\epsilon)^{j_0}} \leq (1+3\epsilon)\chi,$$

and

$$Y_{j_0} > (1-\epsilon')\frac{n_{rest}}{(1+\epsilon)^{j_0}} \geq (1-\epsilon')\frac{\chi}{(1+\epsilon)} \geq \frac{\chi}{(1+\epsilon)^2}.$$

□

Fact 5.7. *Let* $Z = \max_{j \in [\mathfrak{a}]}\{j : \frac{\chi}{(1+\epsilon)^2} < Y_j \leq (1+3\epsilon)\chi\}$ *if at least one such* j *exists and* 0 *otherwise. Then*

$$P(Z > \mathfrak{f}_\chi(n_{rest}) + 2) \leq 1 - \frac{1}{\chi'}.$$

Proof. Let $j_0 = \mathfrak{f}_\chi(n_{rest})$ and consider fixed $j' > j_0 + 2$. We have,

$$E(Y_{j'}) = \frac{n_{rest}}{(1+\epsilon)^{j'}},$$

and by pairwise independency of Xs

$$Var(Y_{j'}) = n_{rest}Var(X_{j',1}) = \frac{n_{rest}}{(1+\epsilon)^{j'}}(1 - \frac{1}{(1+\epsilon)^{j'}}) \leq \frac{n_{rest}}{(1+\epsilon)^{j'}} \leq \frac{\chi}{(1+\epsilon)^{j'-j_0}}.$$

Thus,

$$P(Y_{j'} > \frac{\chi}{(1+\epsilon)^2}) = P(Y_{j'} - E(Y_{j'}) > \frac{\chi}{(1+\epsilon)^2} - E(Y_{j'})) \leq \frac{Var(Y_{j'})}{(\frac{\chi}{(1+\epsilon)^2} - E(Y_{j'}))^2} \leq$$

$$\frac{\chi}{(1+\epsilon)^{j'-j_0}(\frac{\chi}{(1+\epsilon)^2} - \frac{n_{rest}}{(1+\epsilon)^{j'}})^2} \leq \frac{\chi}{(1+\epsilon)^{j'-j_0}(\frac{\chi}{(1+\epsilon)^2} - \frac{\chi}{(1+\epsilon)^{j'-j_0}})^2} =$$

$$\frac{1}{\chi}\frac{1}{(1+\epsilon)^{j'-j_0}(\frac{1}{(1+\epsilon)^2}-\frac{1}{(1+\epsilon)^{j'-j_0}})^2} \le \frac{1}{\chi}\frac{1}{(1+\epsilon)^{j'-j_0}(\frac{1}{(1+\epsilon)^2}-\frac{1}{(1+\epsilon)^3})^2} \le$$

$$\frac{1}{\chi}\frac{1}{(1+\epsilon)^{j'-j_0}\frac{\epsilon^2}{(1+\epsilon)^6}} \le \frac{(1+\epsilon)^3}{\epsilon^2\chi}\frac{1}{(1+\epsilon)^{j'-j_0-3}}.$$

Clearly $Z = j'$ implies $Y_{j'} > \frac{\chi}{(1+\epsilon)^2}$. Thus, and by union bound over all $j' \ge j_0 + 3$, we have that

$$P(Z > j_0 + 2) \le \frac{(1+\epsilon)^3}{\epsilon^2\chi}\sum_{j'=j_0+3}^{\mathfrak{b}}\frac{1}{(1+\epsilon)^{j'-j_0-3}} \le \frac{(1+\epsilon)^3}{\epsilon^2\chi}\frac{(1+\epsilon)}{\epsilon} \le \frac{1}{\chi'}.$$

□

Corollary 5.8. *Let $Y'_i = \sum_{j\in[n_{rest}]}\alpha_{i,j}X_{i,j}$, where $\alpha_{i,j}$ are arbitrary random zero-one variables. For $Z' = \max_{j\in[\mathfrak{a}]}\{j : \frac{\chi}{(1+\epsilon)^2} < Y'_j \le (1+3\epsilon)\chi\}$, it is true that $P(Z' > \mathfrak{f}_\chi(n_{rest}) + 2) \le \frac{1}{\chi'}$.*

Proof. We have for any j

$$P(Z' = j) \le P(Y'_j > \frac{\chi}{(1+\epsilon)^2}) \le P(Y_j > \frac{\chi}{(1+\epsilon)^2}).$$

Thus, we can repeat the arguments from Fact 5.7. □

Fact 5.9. *Let $\zeta = (1+\epsilon)^{1/Q} - 1$, then $\zeta \ge \frac{\epsilon}{2Q}$.*

Proof. If $\zeta < \frac{\epsilon}{2Q}$, then we have

$$(1+\zeta)^Q - 1 = \sum_{i=1}^{Q}\zeta^i\binom{Q}{i} \le \sum_{i=1}^{Q}\zeta^i Q^i < \sum_{i=1}^{Q}\frac{\epsilon^i}{2^i Q^i}Q^i = \sum_{i=1}^{Q}\frac{\epsilon^i}{2^i} \le \epsilon.$$

Thus, it must be the case that $\zeta \ge \frac{\epsilon}{2Q}$. □

5.2.3 The Algorithm and Proof of Lemma 5.5

Algorithm 5.10. $\underline{\mathfrak{G}(D, \mathcal{H}, \epsilon, \delta)}$

1. *Pick random integer q from $0, \ldots, Q-1$.*
2. *For any $j \in [\mathfrak{a}]$ generate pairwise-independent random hash functions $G_j : [n] \to \{0,1\}$ such that for any $i \in [n]$ $P(G_j(i) = 1) = \frac{1}{(1+\epsilon)^j}$.*
3. *In parallel, apply $\mathfrak{Q}_j = \mathfrak{Q}(D, \mathcal{H}, G_j, \mathfrak{c}, \frac{1}{\chi'})$ for all $j = 0, \ldots, \mathfrak{a}$.*
4. *For all $0 \le j \le \mathfrak{a}$ and all $l = -1, \ldots, \mathfrak{b}$ compute $\mathcal{Y}_{l,j}$ that is a number of elements returned by $\mathfrak{Q}_j$ in the range $[(1+\zeta)^q(1+\epsilon)^l, (1+\zeta)^q(1+\epsilon)^{l+1})$.*
5. *For every $l \in [\mathfrak{b}]$ compute $\mathcal{Z}_l = \max_{j>0}\{j : \frac{\chi}{(1+\epsilon)^2} < \mathcal{Y}_{l,j} \le (1+3\epsilon)\chi\}$; define $\mathcal{Z}_l = 0$ if no such j exists.*
6. *Return $(1+\zeta)^q\sum_{l\in[\mathfrak{b}]}(1+\epsilon)^{\mathcal{Z}_l+l}\mathcal{Y}_{l,\mathcal{Z}_l}$.*

Let $\mathcal{F}$ be a fixed function that defines vector $V = \mathcal{F}(D, \mathcal{H})$ with non-negative entries v_i such that $L_\infty(V) = poly(n, m)$. Define q to be a uniform random integer from $0, \ldots, Q-1$. For $l = -1, \ldots, \mathfrak{b}$, define a "layer" S_l as a set of all v_is in the range $[(1+\zeta)^q(1+\epsilon)^l, (1+\zeta)^q(1+\epsilon)^{l+1})$. Denote by s_l the number of elements in S_l. For any l define a *left boundary sub-layer* $S_{l,left}$ as a set of all v_is in the range $[(1+\zeta)^{q-1}(1+\epsilon)^l, (1+\zeta)^q(1+\epsilon)^l)$ and $s_{l,left}$ to be its size. For any l define a *right boundary sub-layer* $S_{l,right}$ as a set of all v_is in the range $[(1+\zeta)^q(1+\epsilon)^l, (1+\zeta)^{q+1}(1+\epsilon)^l)$ and $s_{l,right}$ to be its size. Let $\mathfrak{S}$ be the set of all element in boundary (left or right) sublayers. It is straightforward to see the total weight of the elements in $\mathfrak{S}$ is small, w.h.p.:

Fact 5.11. $P(\sum_{v_i \in \mathfrak{S}} v_i \geq \frac{20}{Q}|V|) \leq 0.1$.

Proof. For a fixed v_i, let $j = Qx + y, 0 \leq y < Q$ be such that $(1+\zeta)^{j-1} < v_i \leq (1+\zeta)^j$. Then, $P(v_i \in \mathfrak{S}) = P(q-1 \leq y \leq q) = \frac{2}{Q}$. Thus, by Markov inequality, $P(\sum_{v_i \in \mathfrak{S}} v_i \geq \frac{20}{Q}|V|) \leq 0.1$. □

Proof of Lemma 5.5 We prove that Algorithm 5.10 satisfies the requirement of the lemma. Let $\mathcal{B}$ be the event that for all j, $\mathfrak{B}_\mathfrak{j}$ returns a $\mathfrak{c}$-cover of VG_j (see Definition 5.1). By parameters of $\mathfrak{B}_\mathfrak{j}$ and by the union bound $P(\mathcal{B}) \geq 1 - \frac{\mathfrak{a}}{\chi'}$ for any fixed functions G_j. Let $\mathcal{D}$ be the event that $\sum_{v_i \in \mathfrak{S}} v_i < \frac{20}{Q}|V|$. By Fact 5.11, we have $P(\mathcal{D}) \geq 0.9$. In the remainder of this section we assume that $\mathcal{B}, \mathcal{D}$ are true. The key observation is that if $\mathcal{B}$ is true then any $v_i \notin \mathfrak{S}$ is not misclassified; i.e., if an approximation of v_i is returned, then it will belong to the same layer as v_i.

I. Upper Bound

To prove the upper bound, we distinguish between *large* and *small* layers. A layer S_l is large if $\tilde{s}_l = s_l + s_{l,left} + s_{l,right} > \chi$, and small otherwise. Consider a fixed l; if S_l is a large layer, then Corollary 5.8 is applicable as follows. Let $X_{i,j}$ be the indicator of the event that $G_j(i) = 1$, and let $v_{i_1}, \ldots, v_{i_{\tilde{s}_l}}$ be the elements from $S_l \cup S_{l,left} \cup S_{l,right}$. Let $\alpha_{i,j}$ be the indicator random variable that the approximation of v_i will be counted by$\mathcal{Y}_{l,j}$. Since $\mathcal{B}$ is true, no elements outside of $S_l \cup S_{l,left} \cup S_{l,right}$ can be counted. Thus, we can write

$$\mathcal{Y}_{l,j} = \sum_{t=1}^{\tilde{s}_l} X_{i_t,j}\alpha_{i_t,j}$$

and apply Corollary 5.8 with $n_{rest} = \tilde{s}_l$ and an appropriate enumeration of Xs. Therefore, by Corollary 5.8, w.p. at least $1 - \frac{1}{\chi'}$,

$$\mathcal{Z}_l \leq \mathfrak{f}_\chi(\tilde{s}_l) + 2.$$

Consider the case that $\mathcal{Z}_l > 0$. Then, by definition of $\mathcal{Z}_l$, we have $\mathcal{Y}_{l,\mathcal{Z}_l} \leq (1+3\epsilon)\chi$, and thus by definition of $\mathfrak{f}_\chi$:

$$(1+\epsilon)^{\mathcal{Z}_l}\mathcal{Y}_{l,Z_l} \leq (1+\epsilon)^{\mathfrak{f}_\chi(\tilde{s}_l)+2}(1+3\epsilon)\chi \leq (1+\epsilon)^6\tilde{s}_i.$$

Also if $\mathcal{Z}_l = 0$ then $\mathcal{Y}_{l,0} \leq \tilde{s}_l$, assuming $\mathcal{B}$. In this case we have $(1+\epsilon)^{\mathcal{Z}_i}Y_{i,Z_i} \leq \tilde{s}_i$. Thus, for any large layer, we have w.p. at least $1 - \frac{1}{\chi'}$:

$$(1+\epsilon)^{\mathcal{Z}_l}\mathcal{Y}_{l,Z_l} \leq (1+\epsilon)^6\tilde{s}_l.$$

Consider the case when S_l is small. For the purposes of our analysis, we can add to $\mathcal{Y}_{l,j}$ arbitrary elements $v_{\tilde{s}_l+1}, \ldots v_{\chi+1} \notin S_l \cup S_{l,left} \cup S_{l,right}$ and define $\alpha_{i_t,j} \equiv 0$ for all j and for all $t > \tilde{s}_l$. Thus, the above bounds will be valid. Thus, we conclude that for every layer S_l the approximation of its cardinality exceeds $(1+\epsilon)^6\tilde{s}_l$ w.p. at most $\frac{1}{\chi'}$. By union bound and by Fact 5.11, w.p. at least $1 - \frac{\mathfrak{b}}{\chi'}$:

$$(1+\zeta)^q \sum_{l\in[\mathfrak{b}]}(1+\epsilon)^{\mathcal{Z}_l+l}\mathcal{Y}_{l,\mathcal{Z}_l} \leq (1+\zeta)^q \sum_{l\in[\mathfrak{b}]}(1+\epsilon)^{l+6}\tilde{s}_l \leq$$

$$\sum_{i\in[n]} v_i(1+\epsilon)^7 + \sum_{v_i\in\mathfrak{S}} v_i(1+\epsilon)^7 \leq (1+\epsilon)^7(1+20\epsilon)|V|.$$

II. Lower Bound

Now let us prove the lower bound. Assuming $\mathcal{B}$, the only elements from S_l that cannot be counted by $\mathcal{Y}$ are those from $S_{l-1,right}$ and $S_{l+1,left}$. Let $\hat{S}_l = S_l \backslash (S_{l-1,left} \cup S_{l+1,right})$ and let $\hat{s}_l = s_l - s_{l+1,left} - s_{l-1,right}$ to be its size. We change a definition of a *large* layer; S_l is large if $\hat{s}_l > \chi$, and *small* otherwise. Consider an $\frac{\epsilon}{\chi'}$-significant layer $\hat{S}_l$.

II.1. large layers

First, let us assume that $\hat{S}_l$ is large. Let $v_{i_1}, \ldots, v_{i_{\hat{s}_l}}$ be elements from $\hat{S}_l$. Let $X_{i,j} = \mathbf{1}_{G_j(i)=1}$ and let $Y_{l,j} = \sum_{t=1}^{\hat{s}_l} X_{i_t,j}$; i.e., $Y_{l,j}$ is the number of elements among $v_{i_1}, \ldots, v_{i_{\hat{s}_l}}$ that has not been zeroed by G_j. Consider an event $\mathcal{A}$ that $\frac{\chi}{(1+\epsilon)^2} < Y_{l,\mathfrak{f}_\chi(\hat{s}_l)} \leq \chi(1+3\epsilon)$. By Fact 5.6 we have

$$P(\mathcal{A}) \geq 1 - \frac{1}{\chi'}.$$

Let $\tilde{R} = \sum_{v_i \notin \hat{S}_l} G_{\mathfrak{f}_\chi(\hat{s}_l)}(i) v_i$ be the total weight of all elements that do not belong to $\hat{S}_l$ and contribute to $|VG_{\mathfrak{f}_\chi(\hat{s}_l)}|$. We have

$$E(\tilde{R}) = \frac{1}{(1+\epsilon)^{\mathfrak{f}_\chi(\hat{s}_l)}} \sum_{v_i \notin \hat{S}_l} v_i \leq \frac{|V|}{(1+\epsilon)^{\mathfrak{f}_\chi(\hat{s}_l)}}.$$

Consider the event $\mathcal{C}$ that $\tilde{R} \leq \frac{\chi'|V|}{(1+\epsilon)^{\mathfrak{f}_\chi(\hat{s}_l)}}$. We have by Markov inequality that

$$P(\mathcal{C}) \geq 1 - \frac{1}{\chi'}.$$

Below we prove that all elements from $\hat{S}_l$ will belong to $\mathfrak{c}$-cover returned by $\mathfrak{O}_{\mathfrak{f}_\chi(\hat{s}_l)}$. Recall that for any $v_i \in \hat{S}_l$ we have $(1+\zeta)^q(1+\epsilon)^{l-1} < v_i \leq (1+\zeta)^q(1+\epsilon)^l$. Thus, for every $v_i \in \hat{S}_l$ since $\hat{S}_l$ is $\frac{\epsilon}{\chi'}$-significant, $\mathcal{C}$ is true and by definition of $\mathfrak{f}_\chi$:

$$v_i \geq (1+\zeta)^q(1+\epsilon)^{l-1} \geq \frac{\epsilon}{\chi'}\frac{1}{(1+\epsilon)\hat{s}_l}|V| \geq \frac{\epsilon}{\chi'^2}\frac{1}{(1+\epsilon)\hat{s}_l}\tilde{R}(1+\epsilon)^{\mathfrak{f}_\chi(\hat{s}_l)} \geq$$

$$\frac{\epsilon}{2\chi\chi'^2} \sum_{v_{i'} \notin \hat{S}_l} G_{\mathfrak{f}_\chi(\hat{s}_l)}(i) v_{i'}.$$

Since $\mathcal{A}$ is true it follows that $Y_{l,\mathfrak{f}_\chi(\hat{s}_l)} \leq \chi(1+3\epsilon)$. Thus,

$$v_i \geq (1+\zeta)^q(1+\epsilon)^{l-1} \geq \frac{Y_{l,\mathfrak{f}_\chi(\hat{s}_l)}}{4\chi}(1+\zeta)^q(1+\epsilon)^{l-1} =$$

$$\frac{1}{4\chi} \sum_{v_{i'} \in \hat{S}_l} G_{\mathfrak{f}_\chi(\hat{s}_l)}(i)(1+\zeta)^q(1+\epsilon)^{l-1} \geq \frac{1}{8\chi} \sum_{v_{i'} \in \hat{S}_l} G_{\mathfrak{f}_\chi(\hat{s}_l)}(i) v_{i'}.$$

Thus, we conclude that

$$v_i \geq \frac{\epsilon}{4\chi\chi'^2}|VG_{\mathfrak{f}_\chi(\hat{s}_l)}|.$$

But this bound and $\mathcal{B}$ imply that all $v_i \in \hat{S}_l$ with $G_{\mathfrak{f}_\chi(\hat{s}_l)}(i) = 1$ will be found by $\mathfrak{O}_{\mathfrak{f}_\chi(\hat{s}_l)}$ and counted by $\mathcal{Y}_{l,\mathfrak{f}_\chi(\hat{s}_l)}$. Thus

$$\mathcal{Y}_{l,\mathfrak{f}_\chi(\hat{s}_l)} \geq Y_{l,\mathfrak{f}_\chi(\hat{s}_l)} \geq \frac{\chi}{(1+\epsilon)^2}. \tag{5}$$

Let $\mathfrak{O}_l = S_{l,left} \cup S_{l-1,right} \cup S_{l+1,left} \cup S_{l,right} \subseteq \mathfrak{S}$. Let $\mathfrak{o}_l$ be the number of elements in $\mathfrak{O}_l$. Then since $\mathcal{D}$ is true, we have:

$$(1+\zeta)^{q-1}(1+\epsilon)^l\hat{s}_l \geq \sum_{v_i \in \hat{S}_l} v_i \geq \frac{\epsilon}{\chi'}|V| \geq \frac{20}{\epsilon Q}|V| \geq \frac{1}{\epsilon}\sum_{v_i \in \mathfrak{S}} v_i \geq$$

$$\frac{1}{\epsilon}\sum_{v_i \in \mathfrak{O}_l} v_i \geq \frac{1}{\epsilon}\mathfrak{o}_l(1+\zeta)^{q-1}(1+\epsilon)^{l-1}.$$

Thus,

$$\mathfrak{o}_l \leq \epsilon(1+\epsilon)\hat{s}_l \leq 2\epsilon\hat{s}_l.$$

Consider $\check{Y} = \sum_{v_i \in \hat{S}_l \cup \mathfrak{O}_l} G_{\mathfrak{f}_\chi(\hat{s}_l)}(i)$. Assuming $\mathcal{B}$, only elements from $\hat{S}_l \cup \mathfrak{O}_l$ can contribute to $\mathcal{Y}_{l,\mathfrak{f}_\chi(\hat{s}_l)}$, and thus $\mathcal{Y}_{l,\mathfrak{f}_\chi(\hat{s}_l)} \leq \check{Y}$. Further, we have

$$E(\check{Y}) = \frac{\hat{s}_l + \mathfrak{o}_l}{(1+\epsilon)^{\mathfrak{f}_\chi(\hat{s}_l)}} \leq \frac{(1+2\epsilon)\hat{s}_l}{(1+\epsilon)^{\mathfrak{f}_\chi(\hat{s}_l)}} \leq (1+2\epsilon)\chi.$$

Also, by pairwise independence of $G_{\mathfrak{f}_\chi(\hat{s}_l)}$, we have $Var(\check{Y}) \leq E(\check{Y})$. Thus, by Chebyshev inequality:

$$P(\check{Y} > (1+3\epsilon)\chi) = P(\check{Y} - E(\check{Y}) > (1+3\epsilon)\chi - E(\check{Y})) \leq$$

$$P(\check{Y} - E(\check{Y}) \geq \epsilon\chi) \leq \frac{Var(\check{Y})}{\epsilon^2\chi^2} \leq \frac{(1+2\epsilon)}{\epsilon^2\chi} \leq \frac{1}{\chi'}.$$

Therefore,

$$P(\mathcal{Y}_{l,\mathfrak{f}_\chi(\hat{s}_l)} \geq (1+3\epsilon)\chi) \leq \frac{1}{\chi'}. \tag{6}$$

By (5) and (6), w.p. at least $1 - \frac{2}{\chi'}$ we have $\frac{\chi}{(1+\epsilon)^2} \leq \mathcal{Y}_{l,\mathfrak{f}_\chi(\hat{s}_l)} \leq (1+3\epsilon)\chi$, in which case $\mathcal{Z}_l \geq \mathfrak{f}_\chi(\hat{s}_l) > 0$ and thus by definitions of $\mathcal{Z}_l$ and $\mathfrak{f}_\chi$:

$$(1+\epsilon)^{\mathcal{Z}_l}\mathcal{Y}_{l,\mathcal{Z}_l} \geq (1+\epsilon)^{\mathfrak{f}_\chi(\hat{s}_l)}\frac{\chi}{(1+\epsilon)^2} \geq \frac{\hat{s}_l}{(1+\epsilon)^2}.$$

II.2. small layers

Similarly, if $\hat{S}_l$ is small and $\mathcal{Z}_l > 0$ we have

$$(1+\epsilon)^{\mathcal{Z}_l}\mathcal{Y}_{l,Z_l} \geq \mathcal{Y}_{l,Z_l} \geq \frac{\chi}{(1+\epsilon)^2} \geq \frac{\hat{s}_l}{(1+\epsilon)^2}.$$

Otherwise if $Z_l = 0$ then, w.h.p. $Y_{l,0} \geq \hat{s}_l$. Indeed, for every $v_i \in \hat{S}_l$ we have that

$$v_i \geq \frac{\epsilon}{\chi'}\frac{1}{(1+\epsilon)\hat{s}_l}|V| \geq \frac{\epsilon}{(1+\epsilon)\chi\chi'}|V|.$$

Thus, $\mathfrak{Q}_0$ will return approximations of all elements from $\hat{S}_l$ w.p. at least $1 - \frac{1}{\chi'}$; and all approximations will be counted towards $\mathcal{Y}_{l,0}$. Thus $(1+\epsilon)^{\mathcal{Z}_l}\mathcal{Y}_{l,Z_l} \geq \hat{s}_l$.

II.3. putting it all together

By union bound, for all l such that $\hat{S}_l$ is $\frac{\epsilon}{\chi'}$-significant layers, w.p. at least $1 - \frac{\mathfrak{b}}{\chi'}$ we have

$$(1+\epsilon)^{\mathcal{Z}_l}\mathcal{Y}_{l,Z_l} \geq \frac{\hat{s}_l}{(1+\epsilon)^2}.$$

Note that

$$|V| = \sum_{v_i \in \mathfrak{S}} v_i + \sum_{l \in [\mathfrak{b}]} \sum_{v_i \in \hat{S}_l} v_i.$$

Let $\mathcal{L}$ be the set of all l such that $\hat{S}_l$ is $\frac{\epsilon}{\chi'}$-significant. Assuming $\mathcal{D}$, we have for sufficiently large n:

$$\sum_{v_i \in \mathfrak{S}} v_i + \sum_{l \in \bar{\mathcal{L}}} \sum_{v_i \in \hat{S}_l} v_i \leq 20 \frac{\epsilon^2}{\log n} |V| + \mathfrak{b} \frac{\epsilon}{\chi'} |V| \leq \epsilon |V|.$$

We have obtained that w.p. at least $1 - \frac{\mathfrak{b}}{\chi'}$:

$$(1+\zeta)^q \sum_{l \in [\mathfrak{b}]} (1+\epsilon)^{\mathcal{Z}_l + l} \mathcal{Y}_{l, \mathcal{Z}_l} \geq (1+\zeta)^q \sum_{l \in \mathcal{L}} (1+\epsilon)^{\mathcal{Z}_l + l} \mathcal{Y}_{l, \mathcal{Z}_l} \geq \sum_{l \in \mathcal{L}} (1+\zeta)^q (1+\epsilon)^l \frac{\hat{s}_l}{(1+\epsilon)^2} \geq$$

$$\sum_{l \in \mathcal{L}} \sum_{v_i \in \hat{S}_l} \frac{v_i}{(1+\epsilon)^2} \geq \frac{(1-\epsilon)}{(1+\epsilon)^2} |V|.$$

III. Conclusion

We have shown that, w.p. at least $(0.9)(1 - \frac{\mathfrak{a}}{\chi'})(1 - \frac{2\mathfrak{b}}{\chi'}) > 2/3$, the output of Algorithm 5.10 is greater than or equal to $\frac{(1-\epsilon)}{(1+\epsilon)^2}|V|$ and smaller than or equal to $(1+\epsilon)^7(1+20\epsilon)|V|$. By replacing ϵ with an appropriate $\epsilon' = \Omega(\epsilon)$, we obtain an ϵ-approximation of $|V|$.

IV. Memory bounds

We apply $\mathfrak{a}$ algorithms $\mathfrak{Q}$, thus the total memory required for these is $\mathfrak{a}(\mu(n, m, \mathfrak{c}, \frac{1}{\chi'}))$. To generate pairwise-independent functions H, we need $O(\mathfrak{a} \log n)$ memory bits. We also maintain $\mathfrak{a}\mathfrak{b}$ counters $\mathcal{Y}$. In total, by Fact 5.9, we need

$$O\left(\frac{1}{\epsilon} \log(n) \mu(n, m, \frac{\epsilon^7}{\log^3(nm)}, \frac{\epsilon}{\log(nm)}) + \frac{1}{\epsilon^2} \log^2(nm)\right)$$

memory bits. □

6 Proving Lemmas 2.3 and 2.4

Lemma 2.3. *There exists an algorithm $\mathfrak{B}_{k-1}$ that, given a data stream D and an access to hash functions $H_1, \ldots, H_{k-1}$, in one pass obtains an ϵ-approximation of $|T_{k-1}(W(M_{Ind}, H_1, \ldots, H_{k-1}))|$ using memory $O(\frac{1}{\epsilon^2} \log \frac{1}{\delta} \log \frac{nm}{\epsilon\delta})$.*

Proof. For $j \in [n]$, define C_j to be independent random variables with Cauchy distribution. For $\mathbf{i} \in [n]^{k-1}$, denote $\mathcal{H}(\mathbf{i}) = \prod_{l=1}^{k-1} H_l(\mathbf{i}_l)$. Define

$$Z = \sum_{j=1}^{n} C_j \left(\sum_{\mathbf{i} \in [n]^{k-1}} m_{(\mathbf{i},j)} \mathcal{H}(\mathbf{i}) \right)$$

By the arguments from [31], a median of $\Omega(\frac{1}{\epsilon^2} \log \frac{1}{\delta})$ independent Zs is an (ϵ, δ)-approximation of

$$\sum_{j=1}^{n} \left| \sum_{\mathbf{i} \in [n]^{k-1}} m_{(\mathbf{i},j)} \mathcal{H}(\mathbf{i}) \right| = |T_{k-1}(W(M_{Ind}, H_1, \ldots, H_{k-1}))|.$$

To construct Z in a single pass over D, we follow the ideas from [32]. Define $k+1$ random variables $Joint, Margin_1, \ldots, Margin_k$ to be initially equal to 0 and to be updated as follows. Upon receiving a k-tuple $(\mathbf{i}, j) \in [n]^k, \mathbf{i} \in [n]^{k-1}, j \in [n]$, we put $Joint_s = Joint_s + \mathcal{H}(\mathbf{i}) C_j$. For $s < k$, we put $Margin_s = Margin_s + H_s(\mathbf{i}_s)$. Finally we put $Margin_k = Margin_k + C_j$. We have

$$Joint = \sum_{j \in [n]} C_j \sum_{\mathbf{i} \in [n]^{k-1}} f_{(\mathbf{i},j)} \mathcal{H}(\mathbf{i}).$$

Also, for $s < k$ we have

$$Margin_s = \sum_{\mathbf{i}_s \in [n]} f_s(\mathbf{i}_s) H_s(\mathbf{i}_s).$$

Finally

$$Margin_k = \sum_{j \in [n]} f_k(j) C_j.$$

Thus,

$$\prod_{s=1}^{k} Margin_s = \left(\sum_j C_j f_k(j) \right) \left(\sum_{\mathbf{i} \in [n]^{k-1}} \mathcal{H}(\mathbf{i}) \prod_{s=1}^{k} f_s(\mathbf{i}_s) \right) = m^k \sum_j C_j \sum_{\mathbf{i} \in \in [n]^{k-1}} \mathcal{H}(\mathbf{i}) P_{product}((\mathbf{i}, j)).$$

Thus,

$$m^k W - \prod_{s=1}^{k} Margin_s = \sum_{j=1}^{n} C_j \left(\sum_{\mathbf{i} \in [n]^{k-1}} m_{(\mathbf{i},j)} \mathcal{H}(\mathbf{i}) \right).$$

What remains is to analyze the memory bounds. Recall that we don't count the memory of $\mathcal{H}$, which will be analyzed separately. Thus, we need to bound a memory needed to compute Z_s. To compute Z, our algorithm accesses n random variables C_j and computes a sketch that is a weighted sum of C_j. Indyk shows in [31] (see Sections 3.2 and 3.3), that if the coefficients of $C_{j,s}$ are polynomially bounded integers, then it is possible to maintain such a sum with sufficient precision using $O(\log \frac{nm}{\epsilon\delta})$ memory bits. By Fact 3.7,

all entries of $T_s(W(M_{Ind}, \mathcal{H}))$ are polynomially bounded integers; thus, we can repeat the arguments from [31] and the lemma follows.

□

In the reminder of this paper, we assume that $\varpi = O(kn)$. A ϖ-truncated Cauchy variable X is a modified Cauchy variable Y such $X = -\varpi \mathbf{1}_{Y<-\varpi} + Y\mathbf{1}_{-\varpi\leq Y\leq\varpi} + \varpi\mathbf{1}_{Y>\varpi}$.

Definition 6.1. *Let $C_{j,i}, j \in [t], i \in [n]$ be independent random variables where $C_{1,*}$ are Cauchy and $C_{j,*}, j > 1$ are ϖ-truncated Cauchy variables. For every $\mathbf{i} \in [n]^t$ define $C(\mathbf{i}) = \prod_{l=1}^{t} C_{l,\mathbf{i}_l}$. A* product sketch *of t-dimensional tensor M (with entries $m_{\mathbf{i}}, \mathbf{i} \in [n]^t$) is*

$$\mathcal{C}(M) = \sum_{\mathbf{i}\in[n]^t} m_{\mathbf{i}} C(\mathbf{i}).$$

Lemma 6.2. *It is possible to generate in one pass a product sketch of a tensor $T_{s'}(W(M_{Ind}, H_1, \ldots, H_s))$ for any $0 \leq s' \leq s \leq k$.*

Proof. Generate $C_{j,i}, j \in [k - s'], i \in [n]$ random variables as in Definition 6.1. Consider $k + 1$ variables $Joint, Margin_1, \ldots, Margin_k$ initially zero and updated as follows: compute

$$Joint = Joint + \prod_{j\in[s]} H_j(\mathbf{i}_j) \prod_{j\in[k-s']} C_{j,\mathbf{i}_{s'+j}};$$

and for $j \leq s'$

$$Margin_j = Margin_j + H_j(\mathbf{i}_j);$$

and for $j > s$

$$Margin_j = Margin_j + C_{j-s',\mathbf{i}_j};$$

and for $s' < j \leq s$

$$Margin_j = Margin_j + H_j(\mathbf{i}_j)C_{j-s',\mathbf{i}_j}.$$

At the end, we also compute $Product = \prod_{j=1}^{k} Margin_j$. We consider the quantity $m^k Joint - Product$ written in the form $\sum_{\mathbf{i}\in[n]^{k-s'}} C(\mathbf{i})Coef(\mathbf{i})$. Our goal is to compare $Coef(\mathbf{i})$ with the entries of the tensor $T_{s'}(W(M_{Ind}, H_1, \ldots, H_s))$. Let $\mathbf{i} \in [n]^{k-s'}$ be fixed. For $Joint$, a coefficient that corresponds to $C(\mathbf{i})$ is equal to:

$$\sum_{\mathbf{j}\in[n]^{s'}} f_{(\mathbf{j},\mathbf{i})}(\prod_{l=1}^{s'} H_l(\mathbf{j}_l))(\prod_{l=s'+1}^{s} H_l(\mathbf{i}_{l-s'})).$$

For $Product = \prod Margin_j$, a coefficient that corresponds to $C(\mathbf{i})$ is equal to:

$$\sum_{\mathbf{j}\in[n]^{s'}} \left[(\prod_{l=1}^{s'} H_l(\mathbf{j}_l))(\prod_{l=s'+1}^{s} H_l(\mathbf{i}_{l-s'})) \prod_{l=1}^{s'} f_l(\mathbf{j}_l) \prod_{l=s'}^{k} f_l(\mathbf{i}_{l-s'+1}) \right] =$$

$$m^k \sum_{\mathbf{i}\in[n]^{s'}} P_{product}((\mathbf{i},\mathbf{j}))(\prod_{l=1}^{s'} H_l(\mathbf{j}_l))(\prod_{l=s'+1}^{s} H_l(\mathbf{i}_{l-s'})).$$

Thus, the coefficient of $C(\mathbf{i})$ in $m^k Joint - Product$ is

$$\sum_{\mathbf{j}\in[n]^{s'}} m_{(\mathbf{j},\mathbf{i})}(\prod_{l=1}^{s'} H_l(\mathbf{j}_l))(\prod_{l=s'+1}^{s} H_l(\mathbf{i}_{l-s'})).$$

On the other hand, consider $T_{s'}(W(M_{Ind}, H_1, \ldots, H_s))$. The coefficient of $W(M_{Ind}, H_1, \ldots, H_s)$ is $m'_{\mathbf{i}} = m_{\mathbf{i}} \prod_{l=1}^{s} H_l(\mathbf{i}_l)$. Thus, the coefficient of $T_{s'}(W(M_{Ind}, H_1, \ldots, H_s))$ is for $\mathbf{i} \in [n]^{k-s'}$:

$$\sum_{\mathbf{j}\in[n]^{s'}} m'_{(\mathbf{j},\mathbf{i})} = \sum_{\mathbf{j}\in[n]^{s'}} m_{(\mathbf{j},\mathbf{i})}(\prod_{l=1}^{s'} H_l(\mathbf{j}_l))(\prod_{l=s'+1}^{s} H_l(\mathbf{i}_{l-s'})).$$

Thus $m^k Joint - Product$ is the product sketch for $T_{s'}(W(M_{Ind}, H_1, \ldots, H_s))$. It is important to note that the procedure above works for $s' = 0$ as well. □

Fact 6.3. *Let $C_1, \ldots, C_n$ be independent Cauchy variables and let $\alpha_1, \ldots, \alpha_n$ be arbitrary random variables independent of $C_1, \ldots, C_n$. Then*

$$P(|\sum_i C_i\alpha_i| \leq \frac{|\alpha|}{t}) \leq \frac{1}{t}.$$

Proof. By stability, we have $\sum_i C_i\alpha_i \sim C|\alpha|$, where C is a Cauchy variable. Thus,

$$P(|C||\alpha| \leq \frac{1}{t}|\alpha|) \leq \frac{1}{\pi} \int_{-\frac{1}{t}}^{\frac{1}{t}} \frac{1}{1+x^2} \leq \frac{1}{t}.$$

□

Fact 6.4. *Let $\{a_1, \ldots, a_{\mathfrak{n}}\}$ be non-negative real numbers and let $X_i, i \in [\mathfrak{n}]$ be non-negative random variables such that $P(X_i \leq a_i) \leq \frac{1}{\mathfrak{q}}$. Let $X = \sum_{i\in[\mathfrak{n}]} X_i$. Then*

$$P(X \leq \frac{1}{2}\sum_i a_i) \leq \frac{2}{\mathfrak{q}}.$$

Proof. Let $Y_i = a_i$ if $X_i \geq a_i$, and $Y_i = 0$ otherwise. Then

$$E(Y_i) = a_i P(X_i \geq a_i) \geq a_i(1 - \frac{1}{\mathfrak{q}}).$$

Let $Z_i = a_i - Y_i$. Then $Z_i \geq 0$ and $E(Z_i) \leq \frac{a_i}{\mathfrak{q}}$. Let $Y = \sum_i Y_i, Z = \sum_i Z_i$. Then by Markov inequality,

$$P(Z \geq \frac{\mathfrak{q}'}{\mathfrak{q}} \sum_i a_i) \leq \frac{1}{\mathfrak{q}'}.$$

Thus

$$P(\sum_i a_i - Y \geq \frac{\mathfrak{q}'}{\mathfrak{q}} \sum_i a_i) \leq \frac{1}{\mathfrak{q}'}.$$

Thus

$$P(X \leq (1 - \frac{\mathfrak{q}'}{\mathfrak{q}}) \sum_i a_i) \leq P(Y \leq (1 - \frac{\mathfrak{q}'}{\mathfrak{q}}) \sum_i a_i) \leq \frac{1}{\mathfrak{q}'}.$$

Putting $\mathfrak{q}' = \frac{\mathfrak{q}}{2}$, we obtain

$$P(X \leq \frac{1}{2} \sum_i a_i) \leq \frac{2}{\mathfrak{q}}.$$

□

Lemma 6.5. *Let $Y = \sum_{\mathbf{i} \in [n]^k} \prod_{j=1}^k C_{j,\mathbf{i}_j} m_{\mathbf{i}}$ where all C are Cauchy. For any M with entries $m_{\mathbf{i}}$ and for $\mathfrak{q} > 3^k$ we have*

$$P(|Y| \leq \frac{|M|}{(2\mathfrak{q})^k}) \leq \frac{3^k}{\mathfrak{q}}.$$

Proof. We prove the claim by induction on k. For $k = 1$ we have by Fact 6.3:

$$P(|\sum_{l \in [n]} C_{1,l} m_l| \leq \frac{|M|}{2\mathfrak{q}}) \leq \frac{3}{\mathfrak{q}}.$$

Consider $k > 1$. For simplicity of presentation, put $C_l = C_{1,l}$ and

$$Y_l = \sum_{\mathbf{i} \in [n]^{k-1}} \prod_{j=2}^k C_{j,(l,\mathbf{i}_{j-1})} m_{(l,\mathbf{i})}.$$

Then

$$Y = \sum_{l \in [n]} C_l Y_l.$$

We have, by stability of C_ls that, $\sum_{l \in [n]} C_l Y_l \sim C' \sum_{l \in [n]} |Y_l|$ where C' is Cauchy distributed. Thus

$$P(|\sum_{l \in [n]} C_l Y_l| \geq \frac{|M|}{(2\mathfrak{q})^k}) = P(|C'| \sum_{l \in [n]} |Y_l| \geq \frac{|M|}{(2\mathfrak{q})^k}) \geq$$

$$P(|C'|\sum_{l\in[n]}|Y_l| \geq \frac{|M|}{(2\mathfrak{q})^k}, \sum|Y_l| \geq \frac{|M|}{2^k\mathfrak{q}^{k-1}}) \geq$$
$$P(|C'|\sum_{l\in[n]}|Y_l| \geq \frac{\sum|Y_l|}{\mathfrak{q}}, \sum|Y_l| \geq \frac{|M|}{2^k\mathfrak{q}^{k-1}}).$$

We have, by Fact 6.3:

$$P(|C'|\sum|Y_l| \leq \frac{\sum|Y_l|}{\mathfrak{q}}) \leq \frac{1}{\mathfrak{q}}.$$

Denote by M_l the l-th hyperplane of M. By induction for each l:

$$P(|Y_l| \leq \frac{|M_l|}{(2\mathfrak{q})^{k-1}}) \leq \frac{3^{k-1}}{\mathfrak{q}}.$$

Thus, by Fact 6.4:

$$P(\sum|Y_l| \leq \frac{1}{2}\frac{|M|}{(2\mathfrak{q})^{2(k-1)}}) \leq \frac{2*3^{k-1}}{\mathfrak{q}}.$$

By union bound, and since $\frac{1}{\mathfrak{q}} + \frac{2*3^{k-1}}{\mathfrak{q}} \leq \frac{3^k}{\mathfrak{q}}$, the claim is correct. □

Corollary 6.6. *Let $Y = \sum_{\mathbf{i}\in[n]^k}\prod_{j=1}^k C_{j,\mathbf{i}_j}m_{\mathbf{i}}$ where all $C_{j,*}, j > 1$ are ϖ-truncated Cauchy and all $C_{1,*}$ are Cauchy. For any M with entries $m_{\mathbf{i}}$ we have*

$$P(|Y| \leq \frac{|M|}{200^k 3^{k^2}}) \leq \frac{1}{50}.$$

Proof. Consider an event that no Cs is equal to ϖ. Repeating the arguments from [32], the probability that this event does not occur is bounded by

$$\frac{2kn}{\pi}\int_{-\infty}^{-\varpi}\frac{1}{1+x^2} \leq \frac{2kn}{\varpi\pi} \leq \frac{1}{100}.$$

Thus, and by Lemma 6.5:

$$P(|Y| \leq \frac{|M|}{(2\mathfrak{q})^k}) \leq \frac{1}{100} + \frac{1}{100}$$

for $\mathfrak{q} = 100 * 3^k$. □

Lemma 6.7. *Let M be a s-dimensional tensor for $s \leq k$ and let Y be a product sketch of M. I.e.,*

$$Y = \sum_{\mathbf{i}\in[n]^k}\prod_{j=1}^k C_{j,\mathbf{i}_j}m_{\mathbf{i}},$$

where for all $j \in [k], i \in [n]$ the random variables $C_{j,i}$ are independent and $C_{1,*}$ are Cauchy and $C_{j,*}, j > 1$ are truncated Cauchy. Then $|Y|$ is a $\log^k n$-approximation of $|M|$ w.p. at least 0.07.

Proof. We consider $s = k$; the same arguments can be repeated for any $s < k$. Consider $Y_l = |\sum_{\mathbf{i}\in[n]^{k-1}} \prod_{j=2}^{k} C_{j,\mathbf{i}_j} m_{(l,\mathbf{i})}|$, and let $Y' = \sum_{l\in[n]} Y_l$. Indyk [31] shows that for any C with ϖ-truncated Cauchy distribution, it is true that $E(|C|) \leq \log(\varpi^2+1)/\pi + O(1)$. Thus, and by the independency of all Cs, we have:

$$E(Y_l) = E(|\sum_{\mathbf{i}\in[n]^{k-1}} \prod_{j=2}^{k} C_{j,\mathbf{i}_{j-1}} m_{(l,\mathbf{i})}|) \leq \sum_{\mathbf{i}\in[n]^{k-1}} E(|\prod_{j=2}^{k} C_{j,\mathbf{i}_{j-1}}|)|m_{(l,\mathbf{i})}| =$$

$$\sum_{\mathbf{i}\in[n]^{k-1}} \prod_{j=2}^{k} E(|C_{j,\mathbf{i}_{j-1}}|)|m_{(l,\mathbf{i})}| \leq 3\log^{k-1} n \sum_{\mathbf{i}\in[n]^{k-1}} |m_{(l,\mathbf{i})}|.$$

Thus, by Markov inequality:

$$P(|Y'| > 300\log^{k-1} n|M|) \leq \frac{1}{100}.$$

Since $|Y| \leq |Y'|$ the upper bound follows. The lower bound follows from Corollary 6.6 and since for large enough $n, \log n > 200 * 3^k$. □

Lemma 2.4. *There exists an algorithm A_{s_1,s_2} (for any $0 \leq s_2 \leq s_1 \leq k$) that, given a data stream D and an access to hash functions $H_1,\ldots,H_{s_1}$, in one pass obtains a $\log^k n$-approximation of $|T_{s_2}(W(M_{Ind},H_1,\ldots,H_{s_1}))|$ using memory $O(\log(nm)\log\frac{1}{\delta})$.*

Proof. By Lemma 6.2, it is possible to construct a product sketch for $|T_{s_2}(W(M_{Ind},H_1,\ldots,H_{s_1}))|$ in one pass. Also, by Lemma 6.7 the constructed product sketch is a $\log^k n$-approximation of $|T_{s_2}(W(M_{Ind},H_1,\ldots,H_{s_1}))|$ w.p. $\Omega(1)$. Thus, taking a median $O(\log\frac{1}{\delta})$ of independent product sketches results in a $(\log^k n,\delta)$-approximation. It remains to analyze the memory bounds. Repeating the arguments from [32], each product sketch can be constructed with sufficient precision using $O(k\log nm)$ memory bits. Also, the perfectly random variables can be replaced by pseudorandom variables and using the "sorting" argument from [32] (Section 3.2). We repeat the arguments of Indyk and McGregor k times (instead of two as in [32]). □

Acknowledgments

We thank Piotr Indyk for a helpful discussion.

References

[1] C. Aggarwal (editor), "Data Streams: Models and Algorithms," *Springer Verlag*, 2007.

[2] N. Alon, A. Andoni, T. Kaufman, K. Matulef, R. Rubinfeld, N. Xie, "Testing k-wise and almost k-wise independence," *Proceedings of the ACM symposium on Theory of computing*, 2007, pp. 496-505.

[3] N. Alon, N. Duffield, C. Lund, M. Thorup, "Estimating arbitrary subset sums with few probes," *Proceedings of the twenty-fourth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pp. 317–325, 2005.

[4] N. Alon, O. Goldreich, Y. Mansour, "Almost k-wise independence versus k-wise independence," *Inform. Process. Lett.*, 88:107110, 2003.

[5] N. Alon, Y. Matias, M.Szegedy, "The space complexity of approximating the frequency moments". *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pp.20–29, 1996.

[6] B. Babcock, S. Babu, M. Datar, R. Motwani, J. Widom, "Models and issues in data stream systems," *ACM Symposium on Principles of Database Systems*, (2002), pp. 1-16.

[7] Z. Bar-Yossef, T. S. Jayram, R. Kumar, D. Sivakumar, "An Information Statistics Approach to Data Stream and Communication Complexity," *Proceedings of the 43rd Symposium on Foundations of Computer Science*, pp. 209–218, 2002.

[8] Z. Bar-Yossef, R. Kumar, D. Sivakumar, "Reductions in streaming algorithms, with an application to counting triangles in graphs," *ACM-SIAM Symposium on Discrete Algorithms*, 2002, pp. 623-632.

[9] A. Bagchi, A. Chaudhary, D. Eppstein, M. T. Goodrich, "Deterministic sampling and range counting in geometric data streams," *ACM Transactions on Algorithms (TALG)*, v.3 n.2, p.16, May 2007

[10] T. Batu, L. Fortnow, E. Fischer, R. Kumar, R. Rubinfeld, P. White, "Testing random variables for independence and identity," *FOCS*, 2001, pp. 442-451.

[11] T. Batu , L. Fortnow , R. Rubinfeld , W. D. Smith , P. White, "Testing that distributions are close," *Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, p.259, 2000.

[12] T. Batu, R. Kumar, R. Rubinfeld, "Sublinear algorithms for testing monotone and unimodal distributions," *In Proc. 36th Annual ACM Symposium on the Theory of Computing*, pp. 381-390, 2004.

[13] P. Beame, T. S. Jayram, A. Rudra, "Lower bounds for randomized read/write stream algorithms," *In Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing*, pp. 689–698, 2007.

[14] L. Bhuvanagiri, S. Ganguly, D. Kesh, C. Saha, "Simpler algorithm for estimating frequency moments of data streams," *in ACM-SIAM Symposium on Discrete Algorithms*, 2006, pp. 708-713.

[15] V.Braverman, R.Ostrovsky, "Measuring k-Wise Independence of Streaming Data", http://arxiv.org/abs/0806.4790.

[16] A. Chakrabarti, S. Khot, X. Sun, "Nearoptimal lower bounds on the multi-party communication complexity of set disjointness," *in IEEE Conference on Computational Complexity*, 2003, pp. 107-117.

[17] M. Charikar, K. Chen, M. Farach-Colton, "Finding frequent items in data streams," *Theoretical Computer Science*, v.312 n.1, p.3-15, 2004.

[18] M. Charikar, L. OCallaghan, R. Panigrahy, "Better streaming algorithms for clustering problems," *in ACM Symposium on Theory of Computing*, 2003, pp. 30-39.

[19] D. Coppersmith , R. Kumar, "An improved data stream algorithm for frequency moments," *Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms*, pp.151–156 , 2004.

[20] G. Cormode, M. Datar, P. Indyk, S. Muthukrishnan, "Comparing data streams using hamming norms (How to zero in)," *IEEE Trans. Knowl. Data Eng.*, 15 (2003), pp. 529-540.

[21] G. Cormode and S. Muthukrishnan. "An Improved Data Stream Summary: The Count-Min Sketch and its Applications". J. Algorithms, 55(1):5875, April 2005.

[22] G. Cormode, S. Muthukrishnan, "What's New: Finding Significant Differences in Network Data Streams," INFOCOM 2004.

[23] M. Datar, N. Immorlica, P. Indyk, and V.S. Mirrokni. "Locality-Sensitive Hashing Scheme Based on p-Stable Distributions," *Annual Symposium on Computational Geometry (SoCG)*, 2004.

[24] N. G. Duffield, C. Lund, M. Thorup, “Priority sampling for estimation of arbitrary subset sums,” J. ACM 54(6): (2007).

[25] J. Feigenbaum, S. Kannan, A. McGregor, S. Suri, J. Zhang, “Graph distances in the streaming model: the value of space,” *in ACM-SIAM Symposium on Discrete Algorithms*, 2005, pp. 745-754.

[26] J. Feigenbaum, S. Kannan, M. Strauss, M. Viswanathan, “An Approximate L1-Difference Algorithm for Massive Data Streams,” *Proceedings of the 40th Annual Symposium on Foundations of Computer Science*, p.501, 1999.

[27] A. Gal, P. Gopalan. “Lower bounds on streaming algorithms for approximating the length of the longest increasing subsequence,” *In 48th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2007*.

[28] S. Ganguly. “Estimating Frequency Moments of Update Streams using Random Linear Combinations”. *Proceedings of the 8th International Workshop on Randomized Algorithms*, pp. 369-380, 2004.

[29] S. Ganguly, G. Cormode, “On Estimating Frequency Moments of Data Streams,” APPROX-RANDOM 2007.

[30] I. F. Ilyas, V. Markl , P. Haas, P. Brown, A. Aboulnaga, “CORDS: automatic discovery of correlations and soft functional dependencies,” *Proceedings of the 2004 ACM SIGMOD international conference on Management of data*, 2004.

[31] P. Indyk, “Stable distributions, pseudorandom generators, embeddings, and data stream computation,” *J. ACM*, 53 (2006), pp. 307-323.

[32] P. Indyk, A. McGregor, “Declaring Independence via the Sketching of Sketches,” *ACM-SIAM Symposium on Discrete Algorithms*, 2008.

[33] P. Indyk, D. P. Woodruff, “Optimal approximations of the frequency moments of data streams,” *in ACMSymposium on Theory of Computing*, 2005, pp. 202-208.

[34] R. Kimball, Joe Caserta, “The Data Warehouse ETL Toolkit: Practical Techniques for Extracting, Cleanin,” John Wiley & Sons, 2004.

[35] P. Li, "Estimators and tail bounds for dimension reduction in $l_\alpha, (0 \leq \alpha \leq 2)$ using stable random projections," SODA 2008.

[36] P. Li, "Compressed Counting," SODA 2009.

[37] M. Mitzenmacher, S. P. Vadhan, "Why simple hash functions work: exploiting the entropy in a data stream," SODA 2008.

[38] S. Muthukrishnan, "Data Streams: Algorithms And Applications," *Foundations and Trends in Theoretical Computer Science*, Volume 1, Issue 2.

[39] V. Poosala, Y. E. Ioannidis, "Selectivity Estimation Without the Attribute Value Independence Assumption," *Proceedings of the 23rd International Conference on Very Large Data Bases*, pp.486–495, 1997.

[40] R. Rubinfeld, R. A. Servedio, "Testing monotone high-dimensional distributions," *In Proc. 37th Annual ACM Symposium on the Theory of Computing*, pp. 147-156, 2005.

[41] A. Sahai, and Vadhan, S. 1999. Manipulating statistical difference. In Randomization Methods in Algorithm Design (DIMACS Workshop, December 1997), Panos Pardalos, Sanguthevar Rajasekaran, and Jos Rolim, Eds. DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol. 43. American Mathematical Society, Providence, R.I., pp. 251–270.

[42] X. Sun, D. Woodruff, "The Communication and Streaming Complexity of Computing the Longest Common and Increasing Subsequences," *SODA*, 2007.

[43] M. Szegedy, "The DLT priority sampling is essentially optimal," *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pp.150–158, 2006.